Flips & Partitions in Geometric Graphs

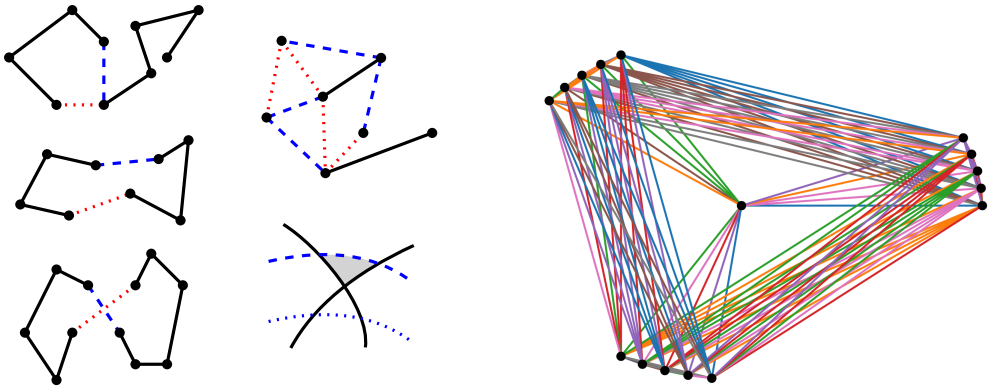

Dissertation zur Erlangung des akademischen Grades

DOKTOR DER NATURWISSENSCHAFTEN

vorgelegt von

JOHANNES OBENAUS

am

Fachbereich Mathematik und Informatik der Freien Universität Berlin

2023

Betreuer und Erstgutachter:	Prof. Dr. Wolfgang Mulzer
Zweitgutachterin:	Prof. Dr. Birgit Vogtenhuber
Datum der Disputation:	10. November 2023

Abstract

This thesis investigates two topics related to fundamental problems in combinatorial geometry. The first being related to *plane graphs*, one of the most widely studied themes in various disciplines related to graph drawing. The second part is concerned with *reconfiguraton problems*, a fundamental field with increasing popularity (see e.g. [BH09, Nis18]).

Edge partitions of complete geometric graphs. The first part of this thesis is concerned with a well-known question posed by Bose, Hurtado, Rivera-Campo, and Wood [BHRW06], who asked whether the edges of every complete geometric graph K_n on an even number of vertices can be partitioned into plane spanning trees. In other words, they asked whether the edges of K_n can be colored in a way such that every color class forms a plane spanning tree. For the special cases that the underlying vertex set is in convex or regular wheel position, a positive answer is known [BHRW06, AHK$^+$17, TCAK19]. However, we prove that the statement is not true in general. Even for partitions into arbitrary plane subgraphs instead of spanning trees we provide a negative answer. Our constructions are based on *bumpy wheel sets* and we give a full characterization which bumpy wheels can be partitioned and which cannot. Additionaly, we provide a characterization for arbitrary wheel sets to admit a partition into plane *double stars* and give a sufficient condition for plane spanning trees.

Finally, we investigate the problem in the broader setting of beyond-planar subgraphs. More precisely, we derive bounds on the number of colors necessary and sufficient to partition a complete geometric graph into k-plane and k-quasi-plane subgraphs. Along the way, we also study the well-known crossing lemma and derive an improvement when restricting to the special case of convex geometric graphs.

Flip graphs. The second part of this thesis is concerned with reconfiguration problems. A natural way to provide structure for a reconfiguration problem is by studying the so-called *flip graph*, which is defined on a ground set $\mathcal{X}$ of objects and a corresponding (local) flip operation. More precisely, the *flip graph* on $\mathcal{X}$ under a given flip operation has a vertex for every element in $\mathcal{X}$ and two vertices are adjacent if and only if the corresponding objects differ by a single flip. For a given ground set and flip operation, an important property one is usually interested in, is whether the flip graph is connected. In the affirmative, more refined questions concerning the diameter, the degree of connectivity, or Hamiltonicity are of interest. We study the following three reconfiguration problems:

- **Flipping plane spanning paths.** For a given point set $S \subset \mathbb{R}^2$ in general position, the ground set $\mathcal{X}$ consists of all plane straight-line paths with vertex set S. The flip operation exchanges a single pair of (potentially crossing) edges. We prove connectedness of the flip graph if the underlying point set S is in wheel position or generalized double circle position. Furthermore, we prove that it suffices to show flip-connectivity for certain subgraphs where the starting edge is fixed.

- **Compatible trees.** For a given simple drawing D of the complete graph K_n, the ground set $\mathcal{X}$ consists of all subdrawings of D that are plane spanning trees. The flip operation exchanges a set of non-crossing edges. We prove connectedness of the flip graph for special classes of drawings, namely cylindrical, monotone, and strongly c-monotone drawings. Furthermore, we prove connectedness of certain subgraphs, corresponding to some classes of graphs, namely stars, double stars, and twin stars.

- **Flipping pseudocircles.** An arrangement of pseudocircles is a finite collection of simple closed curves in the plane such that every pair of curves is either disjoint or intersects in two crossing points. We prove that triangle flips induce a connected flip graph on *intersecting* arrangements, i.e., on arrangements where every pair of pseudocircles intersects. As an intermediate result we also show flip-connectivity on *cylindrical* intersecting arrangements, i.e., arrangements where a single point stabs the interior of every pseudocircle. Moreover, we obtain that the diameter of both flip graphs is cubic in the number of pseudocircles.

Selbstständigkeitserklärung

Ich erkläre gegenüber der Freien Universität Berlin, dass ich die vorliegende Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel angefertigt habe. Die vorliegende Arbeit ist frei von Plagiaten. Alle Ausführungen, die wörtlich oder inhaltlich aus anderen Schriften entnommen sind, habe ich als solche kenntlich gemacht. Diese Dissertation wurde in gleicher oder ähnlicher Form noch in keinem früheren Promotionsverfahren eingereicht. Mit einer Prüfung meiner Arbeit durch ein Plagiatsprüfungsprogramm erkläre ich mich einverstanden.

Datum: _____ Unterschrift: _____
Johannes Obenaus

Acknowledgements

First of all, I would like to express my sincere gratitute to my advisor Prof. Dr. Wolfgang Mulzer for his excellent supervision and guidance.

Furthermore, I thank Birgit Vogtenhuber for reviewing this thesis.

I thank all my co-authors, colleagues, and friends for many fruitful discussions and all their support in so many facettes, making my PhD time a great stage of my life. In particular I want to mention Joachim Orthaber with whom I enjoyed fabulous collaboration.

While preparing this thesis, I also took the exciting journey of raising two wonderful children and I am extremely grateful to everybody who helped with the childcare, my mother, my family, my friends, giving me valuable time to work quietly.

Finally, I am blessed with a wonderful family. Thank you Vivian, Nicolai, and Constantin!

Publications

This thesis is based on the following publications. Note that most of the results in 2. and 3. have been obtained in equal collaboration with Rosna Paul during a research stay of hers (together with other collaborators) and hence, there may be a significant overlap in content with her PhD thesis (this concerns Chapter 7 and Chapter 8).

[AOO+22] Oswin Aichholzer, Johannes Obenaus, Joachim Orthaber, Rosna Paul, Patrick Schnider, Raphael Steiner, Tim Taubner, and Birgit Vogtenhuber. Edge Partitions of Complete Geometric Graphs. *38th International Symposium on Computational Geometry (SoCG)*, 6:1–6:16, 2022.

[AKM+22] Oswin Aichholzer, Kristin Knorr, Wolfgang Mulzer, Nicolas El Maalouly, Johannes Obenaus, Rosna Paul, Meghana M. Reddy, Birgit Vogtenhuber, and Alexandra Weinberger. Compatible Spanning Trees in Simple Drawings of K_n. *30th International Symposium on Graph Drawing and Network Visualization (GD)*, 16–24, 2022.

[AKM+23] Oswin Aichholzer, Kristin Knorr, Wolfgang Mulzer, Johannes Obenaus, Rosna Paul, and Birgit Vogtenhuber. Flipping Plane Spanning Paths. *17th International Conference and Workshops on Algorithms and Computation (WALCOM)*, 49–60, 2023.

[AFO+24] Yan Alves Radtke, Stefan Felsner, Johannes Obenaus, Sandro Roch, Manfred Scheucher, and Birgit Vogtenhuber. Flip Graph Connectivity for Arrangements of Pseudolines and Pseudocircles. *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 4849–4871, 2024.

Contents

Frequently used symbols

$E(G)$	the edge set of a graph G.
$V(G)$	the vertex set of a graph G.
$e(G)$	the number of edges in a graph G.
$f(G)$	the number of faces in a plane graph G.
$K(S)$	the complete straight-line graph on a point set S.
W_k	the regular wheel with $k+1$ vertices.
$BW_{k,\ell}$	the bumpy wheel with k groups and ℓ vertices per group.
$\mathcal{G}_i$	the i-th group in a bumpy wheel.
e^-	in a bumpy wheel, the open halfplane defined by the supporting line through e and not containing the center v_0.
$<_c$	partial order on the non-radial edges (in bumpy wheels): $e <_c f$ if the relative interior of e completely lies in f^-.
$\mathcal{P}(S)$	the set of plane straight-line spanning paths on a point set S.
$\mathcal{T}_D$	the set of plane spanning trees that are a subdrawing of the drawing D.
$\mathcal{F}(\mathcal{X}, f)$	the flip graph on $\mathcal{X}$ under flip operaton f. Also $\mathcal{F}(\mathcal{X})$ is used if the flip operation is clear from the context.

Chapter 1

Introduction

Graph theory shapes almost every part of our modern life, e.g., when routing any type of goods (food, people, electricity, data, etc.) from one place to another. In practice, one is often interested in (efficient) algorithms to solve or approximate a certain problem, such as shortest paths or coloring algorithms. Especially in the light of big data, efficient (graph-theoretic) algorithms are of great importance. One of the aspects that makes graph theory so powerful is the fact that many real-world problems can be phrased as a graph-theoretic problem, usually involving only simple concepts, making it also accessible to non-experts. However, often the theory behind these innocent looking problems is very deep and intricate. Throughout this thesis we will encounter many problems that are simple to phrase but difficult to answer.

When working with graphs, it is very intuitive to visualize these by drawing nodes and edges. These visualizations are not just nice to have, they also play a key role for analyzing the structure of e.g. biological processes, such as protein interactions or networks in social science (see e.g. [JS08, BBD+19]). The field that is concerned with such questions – drawing graphs such that they fulfill certain quality measures – is called *graph drawing*. It is crucial to emphasize that there is a big difference between an *abstract* graph and a *drawing* of this graph, which we will formally define in Chapter 2.

In this thesis, we study two fundamental topics from the field of combinatorial geometry, a field that is concerned with discrete properties of geometric objects, such as points, lines, circles, etc. Classical topics include e.g. packing, covering, partitioning, coloring, crossings, etc.

The first part of this thesis is concerned with the question whether or not the edges of every complete graph drawn with straight-line edges can be colored in a way such that every color class forms a plane spanning tree. The second part is concerned with the topic of reconfiguration (see [Nis18] for a broad introduction of the topic). More specifically, we study connectivity aspects of flip graphs: given a set of objects and a certain flip operation, can any object be transformed into any other by applying suitable flips?

Edge partitions of complete geometric graphs. The focus of this part of the thesis is a well-known and long standing question, generally attributed to Ferran Hurtado and first published in [BHRW06]: is it possible to partition the edges of every complete geometric graph K_n on an even number of vertices into plane spanning trees (see also [OPG])? Or, in other words, can the edges of K_n be colored in a way such that every color class forms a plane spanning tree? For the special cases that the underlying vertex set is in convex or regular wheel position, a positive answer is known [BHRW06, AHK+17, TCAK19]. However, we prove that the statement is not true in general. Even when considering partitions into arbitrary plane subgraphs instead of spanning trees we provide a negative answer. Our constructions are based on *bumpy wheel sets*: intuitively speaking a bumpy wheel is derived from a regular wheel by replacing each extreme point by a group of $\ell \geq 1$ points (see Chapter 2 for the precise definition). We give a full characterization which bumpy wheels can be partitioned and which cannot. In fact, it turns out that there is only one bumpy wheel that can be partitioned into plane subgraphs but not into plane spanning trees. Also for arbitrary wheel sets we provide a characterization to admit a partition into plane *double stars* and give a sufficient condition for plane spanning trees.

Lastly, we investigate the problem in the broader setting of beyond-planar subgraphs (see e.g. [Hon20, DLM19] for further information on beyond planar-graphs). More precisely, we derive bounds on the number of colors needed and sufficient to partition a complete geometric graph into k-plane and k-quasi-plane subgraphs. Along the way, we also study the well-known crossing lemma and derive an improvement when restricting to the special case of convex geometric graphs.

Flip graphs. Reconfiguration is a widely studied topic in discrete mathematics and theoretical computer science [BH09, Nis18]. In many cases, reconfiguration problems can be stated in terms of a *flip graph*, which is defined on a ground set $\mathcal{X}$ of objects and a corresponding (local) flip operation. More precisely, the *flip graph* on $\mathcal{X}$ under a given flip operation has a vertex for every element in $\mathcal{X}$ and two vertices are adjacent if and only if the corresponding objects differ by a single flip. Typically, for a given ground set and flip operation, the first question is whether the flip graph is connected. In the affirmative case, more refined questions regarding diameter, the degree of connectivity, or Hamiltonicity can be of interest. Hamiltonicity of flip graphs is related to Gray codes, cf. the survey by Mütze [Müt22].

A classical example is the flip graph of triangulations of a convex polygon, cf. [Lee89]. The vertex set of this graph are all triangulations of the polygon, and two triangulations are adjacent if one can be obtained from the other by exchanging the common edge of two adjacent triangles by the other diagonal of the convex quadrilateral formed by them. Similar flip graphs have

also been investigated in the context of Delaunay triangulations [Ede01] and for triangulations of general point sets [DLRS10]. Wagner [Wag36] proved connectivity of the flip graph of maximal planar graphs in the 1930s. Its diameter, however, is still not known exactly [BV11]. Further results concerning flip graphs on triangulations, matchings, and trees can be found in [HHN02, HHNR05, Law72, NPTZ20].

In this thesis, we study the following three reconfiguration problems:

- **Flipping plane spanning paths.** For a given point set $S \subset \mathbb{R}^2$ in general position, the ground set $\mathcal{X}$ comprises all plane straight-line spanning paths on S. The flip operation exchanges a single pair of (potentially crossing) edges. We prove connectedness of the flip graph if the underlying point set S is in wheel position or generalized double circle position (which include, e.g., double chains and double circles). Furthermore, we prove that it suffices to show flip-connectivity for certain subgraphs where the starting edge is fixed.

- **Compatible trees.** For a given simple drawing D of the complete graph K_n, the ground set $\mathcal{X}$ comprises the set of all subdrawings of D that are plane spanning trees. The flip operation exchanges a set of non-crossing edges. We prove connectedness of the flip graph for special classes of drawings, namely cylindrical, monotone, and strongly c-monotone drawings. Furthermore, we prove connectedness of certain subgraphs, corresponding to some classes of graphs, namely stars, double stars, and twin stars.

- **Flipping pseudocircles.** An arrangement of pseudocircles is a finite collection of simple closed curves in the plane such that every pair of curves is either disjoint or intersects in two crossing points. We study flip graphs of families of pseudocircle arrangements. We prove that triangle flips induce a connected flip graph on *intersecting* arrangements, i.e., on arrangements where every pair of pseudocircles intersects. First we show that every intersecting arrangement can be flipped into a *cylindrical* arrangement, i.e., an arrangement where a single point stabs the interior of every pseudocircle. Next we flip the cylindrical arrangement into a canonical arrangement, which also shows the connectivity of cylindrical intersecting arrangements of pseudocircles under triangle flips. With a careful analysis we obtain that the diameter of both flip graphs is cubic in the number of pseudocircles. The construction of the two flipping sequences makes essential use of variants of the sweeping lemma for pseudocircle arrangements due to Snoeyink and Hershberger [SH91].

Notation and conventions. Every work in the context of graph drawing requires a certain flexibility between precise distinction of abstract objects and their corresponding drawings on the one hand and concise language on the other hand. Whenever it is clear from the context what object we are refering to, we favour concise language and may use terms such as vertex or point interchangeably.

In Chapter 2 we introduce, in addition to several very basic terms, also the most important results of this thesis in a *formal manner*. To this end, we also define all terms required to state these results. We assume a certain familiarity with the terms introduced in Chapter 2, but may recall some less standard definitions in later chapters for easier readability.

Whenever we consider the indices of a set of k objects in some order (e.g. in cyclic order), we always consider these indices modulo k without further notice. Also, when our indexing begins with 1 rather than 0, we of course consider the representative of the corresponding equivalence class in the range $1, \dots, k$ (rather than $0, \dots, k-1$).

Unless stated otherwise, all graphs, curves, and arrangements in this thesis are simple and all point sets are in general position (we introduce these notions in the following chapter).

Chapter 2

Basic Definitions & Summary of Results

In this chapter, we first introduce basic terms and notation that we will use throughout and then summarize the main results of this thesis in a formal manner.

Graphs and Drawings. A *graph* $G = (V, E)$ consists of a set V of *vertices* and a set E of unordered pairs of vertices, called *edges*. Unless stated otherwise, all graphs in this thesis are *simple*, i.e., do not contain edges of the form $\{v, v\}$ or multiple edges between the same pair of vertices. The *complete graph* K_n on n vertices contains all $\binom{n}{2}$ possible edges. A *drawing* of a graph G is a representation of G in the plane such that every vertex is mapped to a distinct point and every edge is mapped to a simple curve connecting its endpoints and not passing through any other vertex. Two edges sharing a common endpoint are *adjacent*. In a *simple drawing*, every pair of edges has at most one point in common – either a common endpoint or a (proper) *crossing*, i.e., touchings are not allowed. All drawings in this thesis are simple drawings. In a *straight-line drawing* all edges are straight line segments. A graph equipped with a straight-line drawing is called *geometric graph* (as usual, we refer to the drawing and underlying graph interchangeably). In this thesis, we will often consider complete geometric graphs. These graphs are fully determined by the underlying point set. For a point set S in *general position* (i.e. no three points lie on a common line), the complete geometric graph with vertex set S is denoted by $K(S)$.

Isomorphism classes. The question under which conditions two drawings should be considered equal is a bit subtle. Two simple drawings are *strongly isomorphic* if there exists a homeomorphism of the underlying space transforming one drawing into the other. However, often already the set of crossing edge pairs captures all relevant information: Two simple drawings D, D' are *weakly isomorphic* if there exists an isomorphism between the vertex sets such that two edges in D cross if and only if their corresponding edges in D' cross. Essentially, weak isomorphism gives information on *which* edges cross and

strong isomorphism additionally captures the order and orientation of the crossings. In particular, strong isomorphism implies weak isomorphism, but not the other way around.

For complete graphs, weak isomorphism is also determined by the so-called *rotation system*. The *rotation* of a vertex v in a simple drawing is the cyclic order in which the edges incident to v leave the vertex v. The *rotation system* of a simple drawing is the set of rotations of all its vertices. If two simple drawings of a complete graph have the same rotation system, they are weakly isomorphic [PT06]. Conversely, if two simple drawings of a complete graph are weakly isomorphic, they either have the same or inverted rotation systems [Gio22]. For further information on the topic, we refer the reader to the literature, particularly to the work of Kynčl [Kyn09, Kyn11, Kyn13].

Planarity. A graph is called *planar*, if it admits a drawing without crossings. Such a drawing is called *crossing-free* or *plane*.[1] A planar graph equipped with a particular plane drawing is called *plane* graph. It is important to understand that the term "planar" refers to a property of an abstract graph, while "plane" corresponds to a specific drawing.

The notion of planarity can be relaxed in the sense that we allow only a certain number of crossings. We consider two notions of such a relaxation, namely (i) restricting the number of crossings per edge and (ii) restricting the number of pairwise crossing edges (such a set of pairwise crossing edges is called *crossing family*). For an integer $k \geq 0$, we call a graph *k-planar*, if it admits a drawing such that every edge is crossed at most k times. Further, a graph is called *k-quasi-planar* if it admits a drawing with no crossing family of size k. Again, *k-plane* and *k-quasi-plane* drawings/graphs correspond to specific drawings fulfilling the mentioned properties.

For a plane graph G, a *face* Φ of G is a maximal connected region in the complement of the drawing. The edges of G that bound a face Φ are called *bounding edges* and $|\Phi|$ denotes the number of bounding edges, where edges are counted with multiplicities. Furthermore, there is exactly one *unbounded* face and all other faces are *bounded*. The number of faces in a plane graph G is denoted by $f(G)$.

Point sets. Three points in the plane that lie on a common line are *collinear*. Recall that a point set is in general position if there are no three collinear points. All point sets in this thesis are in general position.

A sequence of three non-collinear points p, q, r either forms a left turn or a right turn; in the former we define its *orientation* to be $+1$ and in the latter to be -1. Two point sets $S, S' \subset \mathbb{R}^2$ of equal cardinality are *order-equivalent*

[1] Also the term *embedding* is common in the literature. However, as this term is afflicted with a certain ambiguity whether or not it refers to a *plane* drawing, we do not use it in this thesis.

if there exists a bijection between S and S' that preserves the orientation of all point triples or changes the orientation of every triple. The resulting equivalence classes are called *order types*. Hence, two point sets being of the same order type means that they are order-equivalent. For more information on order types we refer the reader to [PW18].

The *convex hull* $\mathrm{CH}(S)$ of a point set S is the intersection of all convex supersets of S. If there is no danger of confusion, we also use the term convex hull when refering to its boundary or the cycle determining the boundary. Every point of S lying on the boundary of $\mathrm{CH}(S)$ is called *extreme point*. Every point of S that is not an extreme point is called *interior* point.

For geometric graphs, there is a close relation between its rotation system and the order type of the underlying point set: Clearly, the order type determines the rotation system. Conversely, reconstructing the order type from the rotation system is only possible if there are more than three extreme points or the three extreme points are given [ACK+16].

For a straight-line edge $e = pq$, we denote the *supporting line* through p and q by ℓ_e. For a set S of n points, an edge e joining two points of S is called *halving edge* if each of the two open halfplanes determined by ℓ_e contains $\lceil (n-2)/2 \rceil$ or $\lfloor (n-2)/2 \rfloor$ points of S. That is, if n is even there are equally many points on each side of ℓ_e and if n is odd, the numbers of points differ by one. The corresponding line is called *halving line*.

The behaviour of drawings of graphs can be very subtle and difficult to handle. Hence, it is not surprising that many researchers make additional assumptions to restrict the class of drawings to be considered (we will see a couple of examples shortly). In the setting of straight-line drawings, these assumptions usually restrict the configuration of the underlying point set. Also in this thesis, we often consider geometric graphs defined on certain classes of point sets (see also Figure 2.1).

- *convex position* – A set of points is in convex position, if every point is an extreme point.
- *wheel position* – A set of points is in wheel position, if there is exactly one point that is not an extreme point.
- *regular wheel position* – A set of $n+1$ (n odd) points is in regular wheel position, if n points are placed equidistantly along a circle and the remaining point is placed at the circle center. Regular wheels are denoted by W_n. Also note that we require n to be odd in order to assure general position.
- *bumpy wheel position* ([Sch15, Sch16]) – For positive odd integers k and ℓ, the *bumpy wheel* $BW_{k,\ell}$ is derived from the regular wheel W_k by replacing each of the k extreme points by a *group* of ℓ vertices as follows. All vertices (except the center) lie on the convex hull and the vertices within

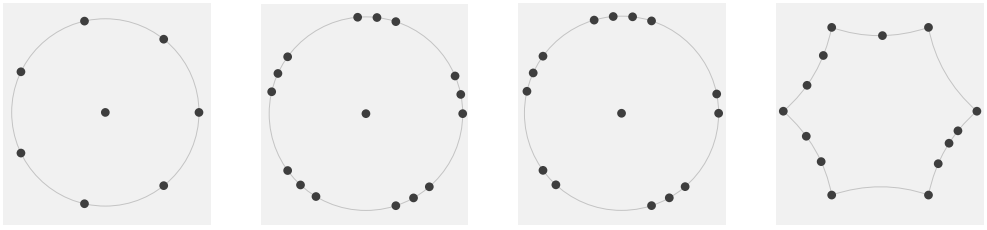

(a) Regular wheel. **(b)** Bumpy wheel. **(c)** Generalized wheel. **(d)** Generalized double circle.

Figure 2.1: Some classes of point sets.

each group are ε-close for some (small enough) $\varepsilon > 0$. In particular, the convex hull of any $\frac{k+1}{2}$ consecutive groups does not contain the center vertex. Note that for $\ell = 1$ we obtain a regular wheel set and for $k = 1$ a point set in convex position and hence, we assume $k, \ell \geq 3$ throughout this thesis.[2]

- *generalized wheel position* – For an odd integer k and a list $N = [n_1, \ldots, n_k]$ of k positive integers, the *generalized wheel* GW_N is derived from the regular wheel W_k by replacing each of the k hull vertices by a group of n_i (for $i = 1, \ldots, k$) vertices in the same way as for bumpy wheels.

  Note that for every wheel set there exists a generalized wheel with the same rotation system (Theorem 5.3).

- *generalized double circle position* – Defining generalized double circles precisely here would be somewhat clumsy as there are several subtleties that become clear later. Therefore, we decided to defer the precise definition to Section 7.3. However, intuitively speaking a generalized double circle is obtained by replacing each edge of the convex hull by a flat enough concave chain of arbitrary size (as depicted in Figure 2.1(d)).

We call a geometric graph whose underlying point set is in convex position *convex graph* or *convex geometric graph*. Similarly, for the other classes of point sets. Furthermore, if there is no danger of confusion we may use the same notation to refer to the point set and the graph.

Special classes of drawings. A drawing of a graph as defined in its most general form allows a lot of freedom concerning the structure of the drawing, which makes those objects difficult handle. However, it is not always necessary to consider drawings in their most general form. For instance, it is well-known (and follows from a simple rerouting argument) that crossing minimal

[2]Strictly speaking the definition requires only k to be odd, but later we need an even number of vertices and hence, we define both, k and ℓ, to be odd right away.

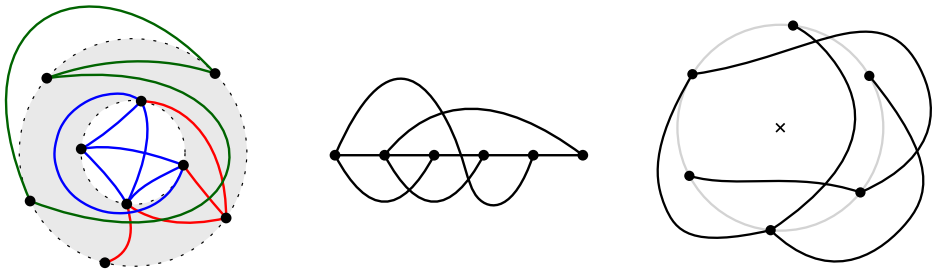

(a) Cylindrical drawing. **(b)** Monotone drawing. **(c)** Strongly c-monotone drawing.

Figure 2.2: Some classes of drawings.

drawings are always simple drawings and hence, whenever concerned with the crossing number, it suffices to consider the class of simple drawings. Unfortunately, also simple drawings provide very little structure and many fundamental questions are unresolved: One of the most prominent open problems is the question whether or not for every $n \in \mathbb{N}$ every simple drawing of K_n contains a plane Hamiltonian cycle. Already in 1988 Rafla [Raf88] conjectured a positive answer; however, despite a lot of effort the question remains open to date. A positive answer is only known for small point sets ($n \leq 9$ [ÁAFM+15]) and special classes of drawings (see Bergold et al. [BFRS23] and Aichholzer et al. [AOV23] for recent results in this context).

Given the intrinsic difficulties of simple drawings, it is natural that many researchers consider drawings with additional structure. One very important class of drawings that we already came across are *straight-line* drawings (geometric graphs), where every edge is drawn as a straight-line segment. We introduce a few more common classes of drawings, relevant for this thesis (see also Figure 2.2):

- *cylindrical drawing* – A drawing, where every vertex is placed on one of two concentric circles and no edge intersects either of the two circles.
- *monotone drawing* – A drawing, where every edge is drawn as x-monotone curve (a curve is *x-monotone* if every vertical line has at most one crossing with the curve).
- *c-monotone drawing* – A drawing, where every vertex is placed on a circle (with center c) and every edge e has the property that every ray emanating from c intersects e at most once.
- *strongly c-monotone drawing* – A c-monotone drawing with the additional property that for every pair of edges e_1, e_2 there exists a ray emanating from c that intersects neither e_1 nor e_2.

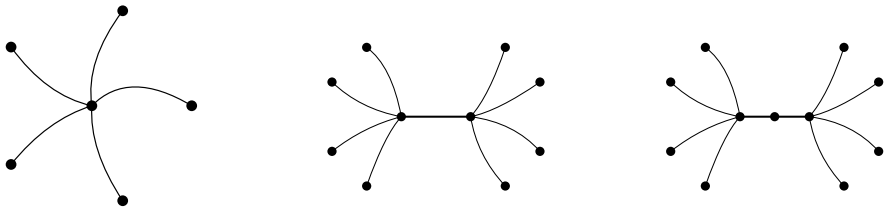

(a) Star. **(b)** Double star. **(c)** Twin star.

Figure 2.3: Some classes of graphs.

We encourage the interested reader to spend a few moments on the nice exercise of verifying Rafla's conjecture concerning the existence of a plane Hamiltonian cycle for special classes of drawings, e.g. for straight-line drawings, cylindrical drawings, or monotone drawings.

Special classes of graphs. Instead of considering special classes of drawings, one can also consider special classes of (abstract) graphs. One such class that we already came across are *complete graphs*, where all $\binom{|V|}{2}$ edges are present. Other very common graph classes are, e.g., *bipartite graphs*, *trees*, and *paths* (we omit the standard definitions, see e.g. [Die17]). A *k-star* is a graph that contains a path $v_0, \ldots, v_k$ of length exactly k such that all remaining vertices have degree one and are adjacent to v_0 or v_k. A 0-star is called *star*, a 1-star is called *double star*, and a 2-star is called *twin star* (see also Figure 2.3).

Partitions. For a set of objects $\mathcal{X}$ and a collection $X_1, \ldots, X_k \subseteq \mathcal{X}$ of subsets of $\mathcal{X}$, we call this collection a

- *covering* of $\mathcal{X}$, if $X_1 \cup \ldots \cup X_k = \mathcal{X}$;
- *packing* of $\mathcal{X}$, if $X_1, \ldots, X_k$ are pairwise disjoint;
- *partition* of $\mathcal{X}$, if $X_1 \cup \ldots \cup X_k = \mathcal{X}$ and $X_1, \ldots, X_k$ are pairwise disjoint.

Applied to the context of partitioning the edge set of a graph, we call a *partition* or *edge partition* of a graph G to be a set of edge-disjoint subgraphs of G whose union is G. We will be interested in partitions of geometric graphs where each subgraph forms a plane spanning tree or just a plane subgraph. Equivalently, one can think of this problem as coloring the edges of a geometric graph in a way such that every color class forms a plane subgraph/spanning tree.

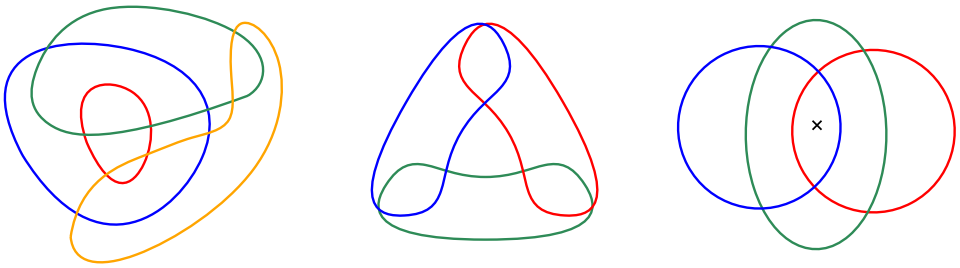

(a) Neither intersecting nor cylindrical. **(b)** Intersecting but not cylindrical. **(c)** Intersecting and cylindrical.

Figure 2.4: Pseudocircle arrrangements.

Arrangements of Pseudocircles. A *pseudocircle* is a simple closed curve C which partitions the plane into a bounded region, the interior $\text{int}(C)$, and an unbounded region, the exterior $\text{ext}(C)$. An *arrangement of pseudocircles* is a finite collection of pseudocircles such that every two pseudocircles either are disjoint or they intersect in exactly one point, where the curves touch, or they intersect in two points, where the curves cross properly. In this thesis, we only consider *simple* arrangements, i.e., there are no touchings and no three curves intersect in a common point. Also note that all curves in this thesis are simple, i.e., do not have self-intersections.

An arrangement of pseudocircles is *intersecting* if every pair of pseudocircles intersects. An intersecting arrangement forms a 4-regular plane graph, possibly with multi-edges. An arrangement of pseudocircles is *cylindrical*[3] if the interiors of all pseudocircles can be stabbed by a single point. See Figure 2.4 for an illustration.

Flip graphs. For a ground set $\mathcal{X}$ of objects, a *flip operation* f transforms an object $x \in \mathcal{X}$ into another object $x' \neq x \in \mathcal{X}$ and vice versa; we say that x and x' differ by a single flip. This symmetry enables us to define an (abstract) graph $\mathcal{F}(\mathcal{X}, f)$ that has vertex set $\mathcal{X}$ and two vertices are adjacent if and only if the corresponding objects differ by a single flip; $\mathcal{F}(\mathcal{X}, f)$ is called *flip graph* on $\mathcal{X}$ under flip operation f. If the flip operation or the ground set is clear from the context, we may omit one or both and just say flip graph $\mathcal{F}$ or flip graph $\mathcal{F}(\mathcal{X})$ on $\mathcal{X}$ for short. For clarity and easier distinguishability we may also introduce a different name.

In order to describe the settings that we investigate for flip connectivity, it suffices to provide the ground set $\mathcal{X}$ and the flip operation f. One of the settings that has most commonly been studied, is the setting where $\mathcal{X}$ denotes

[3] A cylindrical pseudocircle arrangement is not to be confused with a cylindrical drawing. The name "cylindrical" stems from the fact that in both cases the objects can be "nicely" drawn on a cylinder, as we will see later.

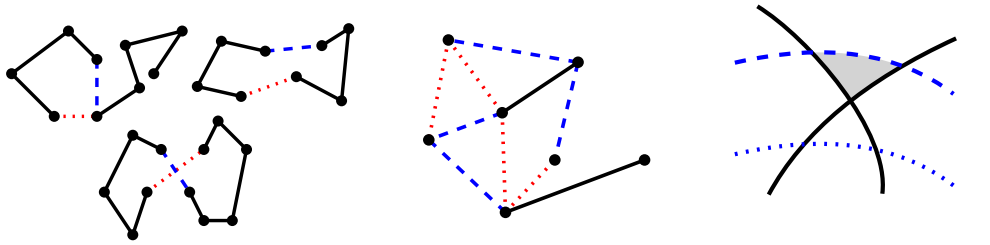

(a) Flipping plane paths. **(b)** Compatible trees. **(c)** Flipping pseudocircles.

Figure 2.5: Illustration of the flip operations in the three settings. In (a) and (b), a set of edges (blue dashed) is replaced by another set of edges (red dotted), while in (c) an arc of a pseudocircle is moved over a crossing of two other pseudocircles.

a class of graphs and the flip operation refers to the exchange of a single pair of edges. In this thesis, we investigate the following three settings (see also Figure 2.5 for an illustration):

- **Flipping plane spanning paths**. Here, $\mathcal{X}$ denotes, for a given point set S, the set of all plane straight-line paths with vertex set S. We denote this set by $\mathcal{P}(S)$. The flip operation exchanges a single pair of (potentially crossing) edges, precisely: $P_1, P_2 \in \mathcal{P}(S)$ differ by a single flip if and only if $E(P_1) \setminus E(P_2) = e_1$ and $E(P_2) \setminus E(P_1) = e_2$, where the edges $e_1 \neq e_2$ may cross.

- **Compatible trees**. Here, $\mathcal{X}$ denotes, for a given simple drawing D of K_n, the set of all subdrawings of D that are plane spanning trees. We denote this set by $\mathcal{T}_D$. Two vertices in the flip graph are adjacent if and only if the corresponding trees are compatible, where two plane graphs are said to be *compatible* if their union is still plane.

  For clarity, we use the term *compatibility graph* instead of flip graph here.

- **Flipping pseudocircles**. Here, $\mathcal{X}$ denotes a class of pseudocircle arrangements with a fixed number of pseudocircles (e.g. cylindrical, intersecting arrangements). The flip operation reroutes a single pseudocircle over a crossing of two others; called a *triangle flip*.

We remark that transformations via compatible trees (the second setting) can be simulated in terms of *crossing free* edge flips as follows. For two compatible trees T_1, T_2, successively perform the following flips: add an edge from $T_2 \setminus T_1$ to T_1 and from the resulting cycle remove an edge that is not in T_2.

2.1 Summary of results

This section is meant as a listing of the most important results of this thesis in a formal manner. In particular, it is not the goal of this section to motivate the results, which we did earlier (and will do in the upcoming chapters). Additionaly, the chapters where to find the results are specified.

Edge partitions of geometric graphs

First, we characterize for which parameters k and ℓ, bumpy wheels can and in particular cannot be partitioned into plane spanning trees or plane subgraphs.

Theorem 2.1. *For odd parameters $k, \ell \geq 3$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane spanning trees if and only if $\ell > 3$.* Chapter 3

Theorem 2.2. *For odd parameters $k, \ell \geq 3$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane subgraphs if and only if $\ell > 5$ or ($\ell = 5$ and $k > 3$).* Chapter 4

Second, we give a necessary condition for generalized wheels to admit a partition into plane spanning trees and characterize which generalized wheels can be partitioned into plane double stars. Recall that generalized wheels always have an odd number of groups.

Theorem 2.3. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning trees if each family of $\frac{k-1}{2}$ consecutive groups contains (strictly) less than $n-2$ vertices.* Chapter 5

Theorem 2.4. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning double stars if and only if there are three families of $\frac{k-1}{2}$ consecutive groups, such that (i) every family contains at most $n-2$ vertices and (ii) every group is in at least one family.* Chapter 5

Finally, we study partitions into beyond-planar substructures, namely k-plane and k-quasiplane subgraphs and spanning trees. Along the way, we also study the well-known crossing lemma and derive an improvement when restricting to the special case of convex geometric graphs (to the best of our knowledge, this special case has not been studied before).

Theorem 2.5. *For a point set S in convex position with $|S| = n \geq 5$, $K(S)$ can be partitioned into $\lceil \frac{n}{3} \rceil$ many 1-plane subgraphs and $\lceil \frac{n}{3} \rceil$ subgraphs are required in every 1-plane partition.* Chapter 6

Theorem 2.6. *For an n-point set S in convex position and every $k \in \mathbb{N}$, $K(S)$ admits a partition into at most $\frac{n}{\sqrt{2k}}$ k-plane subgraphs. More precisely, for every $s > 2$, $K(S)$ admits a $\frac{(s-1)(s-2)}{2}$-plane partition into $\lceil \frac{n}{s} \rceil$ subgraphs. Conversely, for every $k \in \mathbb{N}$, at least $\frac{n-1}{4.93\sqrt{k}}$ subgraphs are required in any k-plane partition of $K(S)$.* Chapter 6

Chapter 6

Lemma 2.7 (convex crossing lemma)**.** *Let G be a graph with n vertices and e edges such that $e \geq \frac{9}{2}n$. Then every straight-line drawing of G in which the vertices of G are placed in convex position has at least*

$$\frac{20}{243}\frac{e^3}{n^2} \approx 0.0823\frac{e^3}{n^2}$$

crossings.

Chapter 6

Theorem 2.8. *Let S be a point set of size $2n$, then the complete geometric graph $K(S)$ can be partitioned into n 3-quasi-plane spanning trees.*

Chapter 6

Theorem 2.9. *Let S be a set of n points in general position and denote the size of a largest crossing family on S by m. Also let $k \in \mathbb{N}$ such that $3 \leq k \leq m$. Then, at least $\lceil \frac{m}{k-1} \rceil$ subgraphs are required and at most $\lceil \frac{m}{k-1} \rceil + \lceil \frac{n-2m}{k-1} \rceil$ subgraphs are needed to partition the complete geometric graph $K(S)$ into k-quasi-plane subgraphs.*

Flip graphs

Recall that $\mathcal{P}(S)$ denotes the set of all plane spanning paths on a point set S. Furthermore, for $p, q \in S$, let $\mathcal{P}(S, p, q)$ be the set of all plane spanning paths for S that start at p and continue with q. In the context of flipping plane spanning paths, we first show that it suffices to prove flip-connectivity on $\mathcal{P}(S, p, q)$ and then prove connectedness of the flip graph for special classes of point sets.

Chapter 7

Theorem 2.10. *Let S be a point set in general position. If, for every $p, q \in S$, the flip graph on $\mathcal{P}(S, p, q)$ is connected, then the flip graph on $\mathcal{P}(S)$ is connected.*

Chapter 7

Theorem 2.11. *Let S be a set of n points in wheel configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter at most $2n - 1$.*

Chapter 7

Theorem 2.12. *Let S be a set of n points in generalized double circle configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter $O(n^2)$.*

Next, we show flip-connectivity in the context of plane spanning trees for broader classes of drawings, but allowing a stronger notion of flips, namely compatibility.

Chapter 8

Theorem 2.13. *Let D be a cylindrical, monotone, or strongly c-monotone drawing of the complete graph K_n. Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected.*

Theorem 2.14. *Let D be a simple drawing of the complete graph K_n and let $\mathcal{T}_D^*$ be the set of all plane spanning stars, double stars, and twin stars on D. Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D^*)$ is connected.* Chapter 8

Last but not least, we consider arrangements of pseudocircles and prove flip-connectivity for two important classes, namely intersecting arrangements and cylindrical intersecting arrangements. All flips are triangle flips.

Theorem 2.15. *The flip graph of arrangements of n pairwise intersecting pseudocircles is connected.* Chapter 9

Theorem 2.16. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles is connected.* Chapter 9

Furthermore, we provide (asymptotically) tight bounds on the diameters of the corresponding flip graphs.

Proposition 2.17. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles has diameter at least $2\binom{n}{3}$ and at most $4\binom{n}{3}$.* Chapter 9

Proposition 2.18. *The flip graph of arrangements of n pairwise intersecting pseudocircles has diameter $\Theta(n^3)$.* Chapter 9

Part I

Edge Partitions of Geometric Graphs

Chapter 3

Partitions into Plane Spanning Trees

The following long-standing open question is the focus of this chapter:

Question 3.1 ([BHRW06])**.** *Can every complete geometric graph on* $2n$ *vertices be partitioned into* n *plane spanning trees?*

Recall that the complete graph K_n contains $\frac{n}{2}(n-1)$ edges and a spanning tree contains $(n-1)$ edges. Hence, only for an even number of vertices the complete graph has the right number of edges to admit a partition into spanning trees. Therefore, unless stated otherwise, we consider complete geometric to have $2n$ vertices in the following.

Related work. Several attempts have been made to answer Question 3.1. When the underlying point set S is in convex position it follows from a result of Bernhart and Kainen [BK79] that $K(S)$ can be partitioned into plane spanning paths, implying a positive answer. Further, Bose et al. [BHRW06] gave a complete characterization of all possible partitions into plane spanning trees for convex point sets. Similarly, when S is a regular wheel set, Aichholzer et al. [AHK$^+$17] showed how to partition $K(S)$ into plane spanning double stars and Trao et al. [TCAK19] recently characterized the structure of all possible partitions into arbitrary plane spanning trees. Furthermore, Aichholzer et al. [AHK$^+$17] provided a positive answer to Question 3.1 for all point sets of (even) cardinality at most 10, obtained by exhaustive enumeration.

Relaxing the requirement that the trees need to be spanning, Bose et al. [BHRW06] showed that if for a point set S in general position, there exists an arrangement of k lines in which every cell contains at least one point from S, then the complete geometric graph on S admits a partition into $2n-k$ plane trees, k of which are plane double stars. This result implies that Question 3.1 has a positive answer if S contains n pairwise crossing segments, which is the case if and only if S has exactly n halving lines [PS99].

For the related *packing* problem where not all edges of the underlying graphs must be covered, Biniaz and García [BG20] showed that $\lfloor n/3 \rfloor$ plane spanning trees can be packed in any complete geometric graph on n vertices, which is currently the best lower bound.

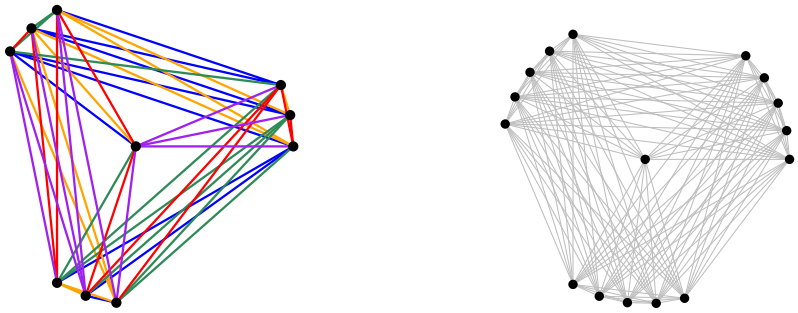

(a) (b)

Figure 3.1: (a) A partition of $BW_{3,3}$ into $n = 5$ plane spanning trees. (b) The bumpy wheel $BW_{3,5}$, which cannot be partitioned into plane spanning trees.

Contribution. We provide a negative answer to Question 3.1 (refuting the prevalent conjecture). For this construction we use bumpy wheel sets. Recall that the bumpy wheel $BW_{k,\ell}$ is obtained from a regular wheel by replacing each of the k extreme points by a group of ℓ points (cf. Chapter 2 for the precise definition).

Our motivation to study bumpy wheels stemmed from the fact that Schnider [Sch15] showed that $BW_{3,3}$ cannot be partitioned into plane double stars. In contrast, this is always possible for complete geometric graphs on regular wheel sets [AHK+17], as well as complete geometric graphs on point sets admitting n pairwise crossing edges [BHRW06] (which also includes convex point sets).

We not only answer Question 3.1, but fully characterize for which parameters k and ℓ, bumpy wheels can be partitioned into plane spanning trees (see Figure 3.1 for an illustration):

Theorem 2.1. *For odd parameters $k, \ell \geq 3$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane spanning trees if and only if $\ell > 3$.*

Organization of this chapter. In Section 3.1 we show how to model this partition problem as an integer linear program (ILP) and provide an implementation that computes solutions for point sets up to roughly 25 points. None of the proofs in this thesis rely on the computer assisted ILP, but it served as a great source of inspiration. In Section 3.2 we show the "if" direction of Theorem 2.1, proving the non-existence of partitions. Whereas in Section 3.3 we prove the "only if" direction of Theorem 2.1, showing the existence of partitions.

3.1 The ILP model

Given a geometric graph $G = (P, E)$ and a fixed number m of available colors as input, our ILP contains a binary variable $x_{e,c} \in \{0, 1\}$ for each edge-color combination, that is, in our setting there are $\binom{2n}{2} \cdot m$ variables. A variable $x_{e,c}$ being 1 then corresponds to edge e receiving color c.

We implement[1] the following constraints, enforcing that every edge receives exactly one color (3.1), crossing edges receive different colors (3.2), ensuring $2n - 1$ edges in each color class (3.3), and forbidding monochromatic triangles (3.4). Clearly, these are necessary but not sufficient constraints:

$$\sum_{c=1}^{m} x_{e,c} = 1 \qquad \forall e \in E \tag{3.1}$$

$$x_{e,c} + x_{f,c} \leq 1 \qquad \forall c \in \{1, \ldots, m\};\ \forall e, f \text{ crossing} \tag{3.2}$$

$$\sum_{e \in E} x_{e,c} = 2n - 1 \qquad \forall c \in \{1, \ldots, m\} \tag{3.3}$$

$$x_{e,c} + x_{f,c} + x_{g,c} \leq 2 \qquad \text{for each triangle } e, f, g;\ \forall c \in \{1, \ldots, m\} \tag{3.4}$$

For $BW_{3,5}$ and $m = 8$ as input, using an industry strength ILP solver, the ILP turns out to be infeasible (taking less than a minute)[2]. Furthermore, for $BW_{3,7}$ and $m = 11$ as input our program reports an infeasible ILP even when omitting the constraints (3.3) and (3.4) (taking roughly 5h). Figure 3.2 shows a partition of $BW_{3,5}$ into plane subgraphs found by the program, when omitting the triangle constraint (3.4).

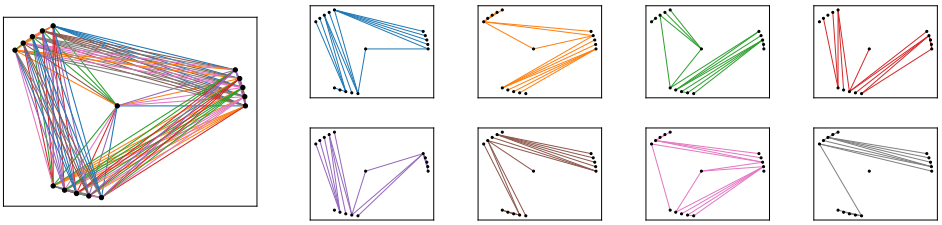

Figure 3.2: Partition of the bumpy wheel $BW_{3,5}$ into 8 plane subgraphs. A partition into plane spanning trees is not possible.

[1] See Appendix A for details concerning the source code.

[2] This and all other experiments were run on an Intel Core i5, 1.6 GHz, 16 GB RAM running macOS Big Sur Version 11.4. All algorithms were implemented in Python 3.9.1, and for solving the ILP we used Gurobi Optimizer Version 9.1.2 with default settings.

3.2 Bumpy wheels that cannot be partitioned into PSTs

In this section, we prove Theorem 2.1. We remark that Theorem 2.1 and Theorem 2.2 are somewhat related, where Theorem 2.2 considers a broader setting. The non-existence direction in Theorem 2.1 follows almost from the non-existence in Theorem 2.2 (the only case that is not covered is $BW_{3,5}$). However, we believe it is more instructive to first consider the more restrictive setting for spanning trees and then extend the results to arbitrary subgraphs. We split the proof of Theorem 2.1 into two parts, starting with the non-existence:

Theorem 3.2. *For any odd parameters $k \geq 3$ and $\ell \geq 5$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane spanning trees.*

We need a few more definitions. For a geometric graph in (bumpy) wheel configuration we denote the center vertex by v_0 and the remaining vertices by $v_1, \ldots, v_{2n-1}$ in clockwise order. We also enumerate the groups in clockwise order: for $i \in \{1, \ldots, k\}$, $\mathcal{G}_i$ denotes the i'th group ($\mathcal{G}_1$ contains v_1, $\mathcal{G}_k$ contains v_{2n-1}). An edge having v_0 as an endpoint is called a *radial edge*, an edge on the convex hull is called a *boundary edge* and all other edges are called *diagonal edges*. For a non-radial edge e, we define e^- to be the open halfplane defined by the supporting line through e and not containing v_0, and similarly e^+ to be the open halfplane containing v_0.

Additionally, we define a partial order $<_c$ on the set of non-radial edges, where $e <_c f$ if the relative interior of e completely lies in f^- (that is, f is "closer" to the center vertex v_0 than e)[3]. Two non-radial edges e, f are *incomparable* with respect to $<_c$, if neither $e <_c f$ nor $f <_c e$ holds (we omit "with respect to $<_c$" if it is clear from the context). In the following, when speaking of an edge e lying in f^- or in f^+ for another edge f, we always refer to the relative interior of e (that is, an endpoint of e may lie on the line through f — which actually means it coincides with an endpoint of f). A non-radial edge e is *maximal* in some set of edges E, if there is no other edge $e' \in E$ such that $e <_c e'$. In the following we often consider maximal diagonal edges of plane spanning trees. *Minimal* edges are defined similarly. See Figure 3.3 for an illustration. Let us emphasize that we never use $<_c$ for radial edges.

Towards the proof of Theorem 3.2, we will first prove several structural results concerning the number and arrangement of radial and diagonal edges in the spanning trees of a potential partition (some of which have a similar flavor as those by Trao et al. [TCAK19]). We show that radial edges must lie between maximal diagonal edges and those maximal diagonal edges need to fulfill certain constraints. We will show that these cannot be satisfied if $\ell \geq 5$.

[3]Note that we defined the relation $e <_c f$ this way around, since e is the (combinatorially) shorter edge.

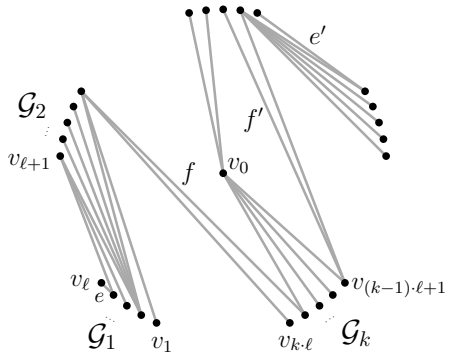

Figure 3.3: Example of a plane spanning tree on the bumpy wheel set $BW_{5,5}$. The diagonal edges f and f' are maximal. The edges e and e' are boundary edges (they are also the only minimal edges).

The following observation follows immediately from convexity and the definition of the partial order $<_c$.

Observation 3.3. *For two non-radial, non-crossing, incomparable edges e, f the vertex sets in e^- and f^- are disjoint and neither e^- nor f^- contains an endpoint of the other edge.*

Note that e and f in the above observation may share an endpoint. Furthermore, for any set of edges E, two maximal edges $e, e' \in E$ are always incomparable.

Lemma 3.4. *Let T be a plane spanning tree of $BW_{k,\ell}$. Then the following properties hold:*

(i) for any diagonal edge $e \in E(T)$, T contains at least one boundary edge in e^-,

(ii) for any pair of incomparable diagonal edges $e, f \in E(T)$, the boundary edges of T in e^- and f^- are distinct, and

(iii) if T contains exactly one maximal diagonal edge, T contains all radial edges of at least $\frac{k-1}{2}$ consecutive groups and at least one more consecutive radial edge (in particular, T contains at least $(\frac{k-1}{2}\ell + 1)$ consecutive radial edges).

Proof. For part (i), let f be a minimal edge of T in e^-. If f^- does not contain any vertex of the input point set, it is a boundary edge and we are done. Otherwise, since T is connected and plane, at least one vertex in f^- has to be connected to an endpoint of f, forming a smaller edge, which is a contradiction to the minimality of f.

Part (ii) follows immediately from Observation 3.3 (distinctness) and part (i) (existence).

Concerning part (iii), let f be the maximal diagonal edge of T. Due to the regularity of the group placement in bumpy wheels, f^+ contains the vertices of at least $\frac{k-1}{2}$ consecutive groups. Hence, since f is the only maximal diagonal edge, all vertices in f^+ need to be reached by radial edges (plus one to connect to f). □

Note that any spanning tree in a partition of $BW_{k,\ell}$ contains a maximal diagonal edge, since the star around v_0 cannot be used in such a partition (as it would not leave any radial edge to reach v_0 for the other trees).

Before diving into further technical details, we first outline the proof from a high level view.

High level proof strategy. Given a bumpy wheel $BW_{k,\ell}$ on $2n$ vertices, assume that a partition into n plane spanning trees $T_0, \dots, T_{n-1}$ exists. We will successively derive several properties concerning the structure of the T_i's, e.g., we show that one tree, say T_0, contains many radial edges. Putting all structural results from Proposition 3.5 to Proposition 3.9 together we will see that these cannot be satisfied if $\ell \geq 5$.

Proposition 3.5. *Let $T_0, \dots, T_{n-1}$ be a partition of $BW_{k,\ell}$ into plane spanning trees (if it exists). Then exactly one of those trees, say T_0, contains a single boundary edge and a single maximal diagonal edge and all other $n-1$ trees contain exactly two boundary edges and exactly two maximal diagonal edges each. In particular, any diagonal edge $e \in E(T_i)$ contains exactly one boundary edge of T_i in e^-.*

Proof. Every T_i contains at least one maximal diagonal edge and hence, by Lemma 3.4 (i), also at least one boundary edge.

Since there are $2n-1$ boundary edges in total, at least one tree (w.l.o.g. T_0) contains exactly one boundary edge. By Lemma 3.4 (i) and (ii) it also contains exactly one maximal diagonal edge.

Now, if there was a second spanning tree T_1 in the partition with exactly one maximal diagonal edge, T_0 and T_1 together would use at least $(k-1)\cdot\ell+2$ radial edges (by Lemma 3.4 (iii)). This leaves at most $\ell-2$ radial edges for the remaining $n-2$ trees; clearly not enough (since $n = \frac{k}{2}\ell + \frac{1}{2} > \ell$ for $k \geq 3$).

Hence, all other $n-1$ spanning trees have to contain at least two maximal diagonal edges and therefore at least two boundary edges. However, since we have $2n-1$ boundary edges in total and only one tree contains a single boundary edge, all other $n-1$ trees have to contain exactly two boundary edges. By Lemma 3.4 (ii), they also contain at most, and therefore exactly, two maximal diagonal edges. □

From now on, T_0 always denotes the spanning tree with exactly one boundary edge (when considering a partition into plane spanning trees). Further, we let all radial edges v_0v_i for $i \in \{1, 2, \dots, \frac{k-1}{2}\ell + 1\}$ be part of T_0 (which we can assume without loss of generality due to symmetry).

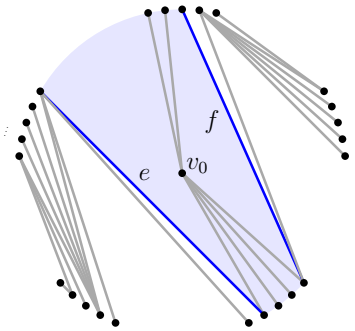

Figure 3.4: Illustration of Lemma 3.6. All edges in the span of the maximal diagonal edges e and f are radial (except e and f itself).

For two non-radial, non-crossing edges e, f, we define the *span* of e and f to be the (closed) region between the two edges, more precisely:

$$\text{span}(e,f) = \begin{cases} \text{cl}(e^+ \cap f^+) & \text{if } e \text{ and } f \text{ are incomparable} \\ \text{cl}(e^+ \cap f^-) & \text{if } e <_c f, \end{cases}$$

where $\text{cl}(\cdot)$ denotes the closure. See also Figure 3.4 for an illustration.

Note, however, that we are more interested in the vertices and edges contained in the span, rather than the region itself. If we want to emphasize this, we may use the notation $V(\text{span}(e,f))$ or $E(\text{span}(e,f))$. In the following we are mostly interested in the span of *maximal* diagonal edges of some plane spanning tree. Figure 3.4 gives an illustration of the following lemma.

Lemma 3.6. *Let $T_0, \dots, T_{n-1}$ be a partition of $BW_{k,\ell}$ into plane spanning trees (if it exists) and let e, f be the maximal diagonal edges of some T_i ($i \neq 0$). Then, all edges of T_i in the span of e and f are radial (except e and f), and all radial edges of T_i lie in the span of e and f.*

Proof. Assume that h is a non-radial edge in the span of e and f. Then, h^- either contains e or f or an additional third boundary edge (or h is a boundary edge itself), a contradiction in any case. Furthermore, any radial edge not contained in $\text{span}(e,f)$ must cross either e or f, and therefore cannot be part of T_i due to planarity. □

We define the *distance* $\text{dist}(e)$ of a non-radial edge e to be the number of vertices in e^- plus one (or in other words, the number of convex hull edges of the underlying point set in $\text{cl}(e^-)$). Clearly, $1 \leq \text{dist}(e) \leq \frac{k+1}{2}\ell - 1$ holds for any non-radial edge e and $\text{dist}(f) < \text{dist}(e)$ holds for any edge $f \subseteq e^-$. It will be convenient to define, for $i \in \{1, \dots, \frac{k+1}{2}\ell - 1\}$:

$$d_i = \frac{k+1}{2}\ell - i. \tag{3.5}$$

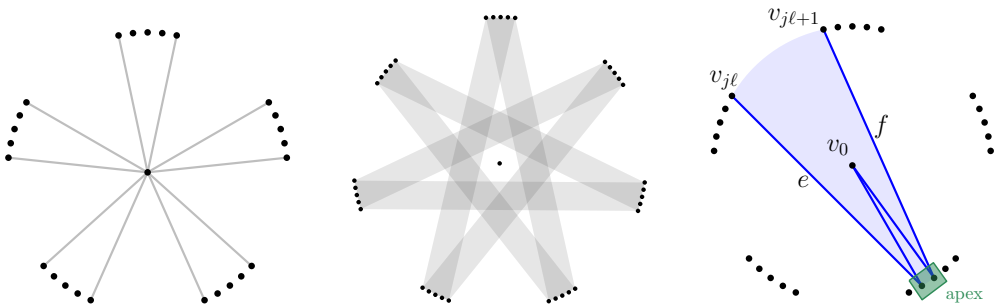

(a) Outmost radial edges. **(b)** Opposite groups. **(c)** Special wedge.

Figure 3.5: Illustration of some terms.

We define it in this (slightly counter-intuitive) way, d_1 being the largest distance, since we are mostly concerned with edges of large distances and thereby aim to improve the readability.

Lemma 3.7. *Consider a plane spanning tree T of a partition of $BW_{k,\ell}$ and let e be a diagonal edge in T of distance* $\operatorname{dist}(e) > 1$. *Then T also contains exactly one of the edges of distance* $(\operatorname{dist}(e) - 1)$ *in* e^-.

Proof. Let f be a maximal edge among the edges of T in e^- and assume $\operatorname{dist}(f) \leq \operatorname{dist}(e) - 2$. Then the span of e and f contains at least one (non-radial) edge $h \notin \{e, f\}$ of T (since either e and f have no endpoint in common and therefore need to be connected via edges in $\operatorname{span}(e, f)$ or $\operatorname{span}(e, f)$ contains at least one vertex v which is neither an endpoint of e nor f and needs to be connected to the rest of T). Since f is maximal in e^-, h and f are incomparable. So, by Lemma 3.4 (ii) this forces at least two distinct boundary edges to be contained in e^-, a contradiction to Proposition 3.5. Therefore, we get $\operatorname{dist}(f) = \operatorname{dist}(e) - 1$.

Furthermore, there are exactly two edges of distance $(\operatorname{dist}(e) - 1)$ in e^-. If $\operatorname{dist}(e) = 2$ these two edges together with e form a cycle and otherwise they cross. Either way, T cannot contain both. □

We need a little more terminology towards the proof of Theorem 3.2 (see also Figure 3.5). We call the first and last vertex of each group *outmost vertices* (and the corresponding radial edges *outmost radial edges*). Note that there are exactly $2k$ outmost radial edges in $BW_{k,\ell}$. Every hull vertex or radial edge that is not outmost, is called an *inside vertex* or an *inside radial edge*.

Furthermore, we define two groups $\mathcal{G}_i, \mathcal{G}_j$ to be *opposite* if $|i - j| = \frac{k-1}{2}$ or $|i - j| = \frac{k+1}{2}$. In particular, each group has two opposite groups and two consecutive groups have exactly one opposite group in common (we call that group *the* opposite group of a pair of consecutive groups).

Let e, f be two maximal (non-crossing) diagonal edges that have an endpoint in a common group. Then the set of vertices of $\operatorname{span}(e, f)$ in the common

group is called *apex*. Note that any apex contains at least one vertex (and this lower bound is attained if the endpoints of e and f coincide).

Moreover, two maximal (non-crossing) diagonal edges $e = \{u, v\}$ and $f = \{u', v'\}$ form a *special wedge* if two endpoints (say u and u') are consecutive outmost vertices of different groups (that is, $u = v_{j\ell}$ and $u' = v_{j\ell+1}$ for some j) and v and v' are inside vertices lying in the opposite group of $\mathcal{G}_j$ and $\mathcal{G}_{j+1}$.

Proposition 3.8. *Let $T_0, \ldots, T_{n-1}$ be a partition of $BW_{k,\ell}$ into plane spanning trees (if it exists) and let T_i ($i \neq 0$) be a spanning tree that does not use any outmost radial edge. Then the two maximal diagonal edges e, f of T_i form a special wedge and T_i uses all radial edges incident to the apex of this wedge.*

Proof. Indeed, by Proposition 3.5, T_i has exactly two maximal diagonal edges. We first argue that all but exactly two radial edges in $\text{span}(e, f)$ must be part of T_i. The subgraph of T_i induced by $V(\text{span}(e, f))$ needs to form a tree. Moreover, the total number of radial edges in $\text{span}(e, f)$ is $|V(\text{span}(e, f))| - 1$. Since T_i uses the two diagonal edges $e, f \in E(\text{span}(e, f))$ and all other edges in the span need to be radial (Lemma 3.6), T_i uses exactly all but two radial edges.

Further note that we cannot have two maximal diagonal edges between the same pair of groups, because diagonal edges between the same pair of groups are either comparable or crossing. Hence, the span of e and f contains at least two outmost vertices, namely in the two distinct groups which contain an endpoint of e and f, respectively. On the other hand, $\text{span}(e, f)$ cannot contain a third outmost vertex, since otherwise T_i has to use an outmost radial edge (by Lemma 3.6 and above argument). Furthermore, by the same argument, none of the two outmost vertices may be in the interior of $\text{span}(e, f)$, i.e., the two outmost vertices are endpoints of e and f. In particular, e and f share a common group and the apex does not contain any outmost vertex. Hence, e and f form a special wedge, as depicted in Figure 3.5.

Moreover, since T_i has to use all but two radial edges in the span, it has to use all radial edges incident to the apex. □

Note that for two spanning trees T_i, T_j ($i \neq j$) not using an outmost radial edge, their apexes are disjoint.

Proposition 3.9. *Let $T_0, \ldots, T_{n-1}$ be a partition of $BW_{k,\ell}$ into plane spanning trees (if it exists). Then for each pair $\mathcal{G}, \mathcal{G}'$ of opposite groups and each $i \in \{1, \ldots, \ell\}$ there is a unique diagonal edge connecting $\mathcal{G}$ and $\mathcal{G}'$ of distance $d_i = \left(\frac{k+1}{2}\ell - i\right)$ that is maximal in its tree.*

Proof. All edges considered in this proof are between $\mathcal{G}$ and $\mathcal{G}'$ without further notice. Observe first that for any $i \in \{1, \ldots, \ell\}$ there are exactly i edges of distance d_i and all edges of the same distance pairwise cross. Also note, for any two edges e, e' with $\text{dist}(e) > \text{dist}(e')$, either $e' \subseteq e^-$ holds or they cross.

In particular, if they do not cross and belong to the same tree, the shorter is not a maximal edge.

Consider now for some $i \in \{2, \dots, \ell\}$ the distance d_i and let $c_1, \dots, c_i$ be the colors[4] used for all i edges of this distance, which are all distinct, since these edges form a crossing family. By Lemma 3.7, the $i-1$ pairwise crossing edges of (the larger) distance d_{i-1} must each use the same color as an edge of distance d_i, w.l.o.g. $c_1, \dots, c_{i-1}$. Hence, the corresponding edges of distance d_i cannot be maximal (again using the argument that there cannot be two maximal diagonal edges between the same pair of groups).

On the other hand, the color c_i cannot be used by any edge of larger distance, since again by Lemma 3.7 there would also be an edge of color c_i with distance d_{i-1}. Hence, indeed the only edge of distance d_i that is maximal in its tree is the one of color c_i.

Lastly, for $i = 1$, the single edge of distance d_1 is clearly maximal. □

Finally, we are ready to prove Theorem 3.2:

Theorem 3.2. *For any odd parameters $k \geq 3$ and $\ell \geq 5$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane spanning trees.*

Proof. Assume to the contrary that there is such a partition $T_0, \dots, T_{n-1}$. There are $2k$ outmost radial edges and T_0 uses (at least) k of them (see Proposition 3.5 and the remark thereafter). Hence, there are at most $k+1$ spanning trees (including T_0) containing an outmost radial edge.

Next, let us count how many spanning trees *without* an outmost radial edge we can have. Since, by Proposition 3.8, the apex of such a tree can neither use any outmost vertex nor any vertex already incident to a radial edge in T_0, there remain $\frac{k+1}{2}(\ell-2)$ possible vertices to be used by apexes, namely the inside vertices of the last $\frac{k+1}{2}$ groups $\mathcal{G}_{\frac{k+1}{2}}, \dots, \mathcal{G}_k$ (the radial edges of those groups not fully used by T_0). Recall that each apex contains at least one vertex.

It is crucial to emphasize that among those last $\frac{k+1}{2}$ groups, group $\mathcal{G}_{\frac{k+1}{2}}$ and group $\mathcal{G}_k$ are opposite (the only opposite pair). Therefore, by Proposition 3.8, two spanning trees with an apex in group $\mathcal{G}_{\frac{k+1}{2}}$ and group $\mathcal{G}_k$, respectively, must each have a maximal diagonal edge between these two groups. Hence, by Proposition 3.9, we can have at most $(\ell-2)$ spanning trees with apex in one of these two groups (instead of $2(\ell-2)$); see Figure 3.6.

In total there are at most $\frac{k-1}{2}(\ell-2)$ spanning trees which do not use an outmost radial edge. Hence, in total we have at most

$$k+1+\frac{k-1}{2}(\ell-2) = \frac{2k+2+k\ell-\ell-2k+2}{2} = \frac{k\ell+1}{2} - \frac{\ell-3}{2} < n$$

spanning trees in our partition, where we used $\ell > 3$ in the last step. This yields the desired contradiction and concludes the proof. □

[4] As mentioned, we associate edges with colors to identify the subgraphs.

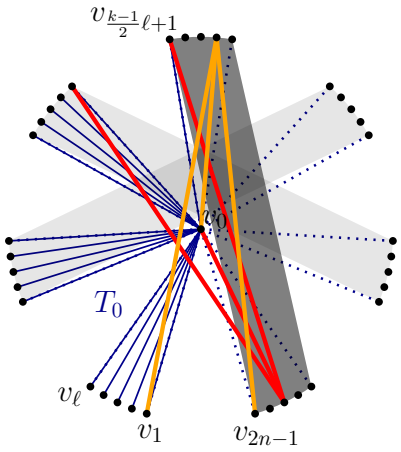

Figure 3.6: In the black stripes (the darker one is the crucial one) the maximal diagonal edges (of those trees without outmost radial edge) need to have distinct distances. That allows $\ell - 2$ many for each stripe. Two spanning trees (red and orange) with apex in group $\mathcal{G}_{\frac{k+1}{2}}$ and group $\mathcal{G}_k$, respectively, both need to have a maximal diagonal edge in the dark stripe.

3.3 Bumpy wheels that can be partitioned into PSTs

The following theorem implies the other direction of Theorem 2.1:

Theorem 3.10. *For any odd $k \geq 3$, there are exactly $4^{k-1} + 4^{k-2}$ non-isomorphic partitions of $BW_{k,3}$ into plane spanning trees.*

Towards the proof of Theorem 3.10, we need a little more terminology. A *partial partition* of a graph G is a set of edge-disjoint subgraphs of G whose union is a subgraph of G (i.e. the edge sets of the subgraphs form a packing of $E(G)$). A subgraph of G, which can be extended to a spanning tree of G, is called *partial tree*. Our construction consists of three steps, where we are starting from a partial partition that fulfills certain properties and show how to extend it to a proper partition step by step adding edges to the subgraphs. Whenever we use the phrase "an edge is covered", we mean that this edge already belongs to a subgraph (of the partial partition).

A partial partition of $BW_{k,\ell}$ into n subgraphs is called *base partition* if the following properties are fulfilled:

(a) all radial edges are covered,

(b) for each pair of opposite groups, there is exactly one diagonal edge of each distance $d_1, \ldots, d_\ell$ covered, and no further diagonal edges are covered,

(c) each of the covered diagonal edges is maximal in its partial tree, and

(d) each partial tree is connected, plane, and non-empty.

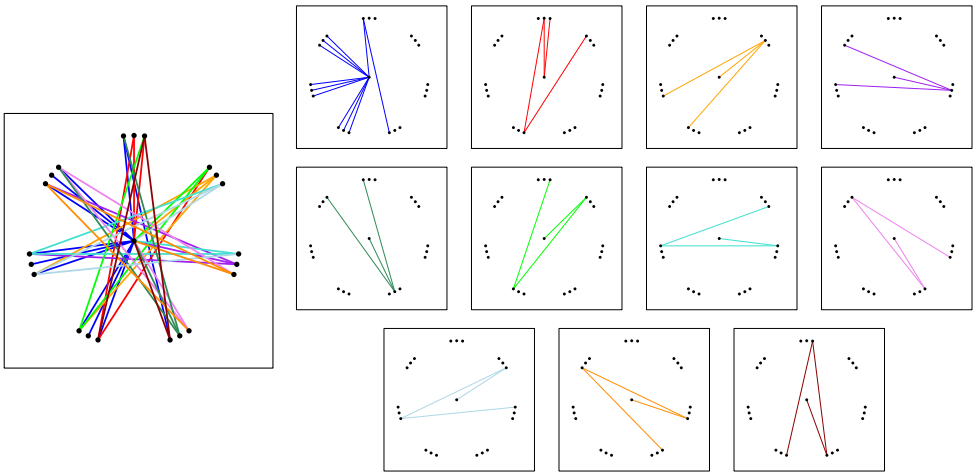

Figure 3.7: An example of a base partition for $BW_{7,3}$.

Figure 3.7 gives an illustration of a base partition. Proving the exact value of $4^{k-1} + 4^{k-2}$ in Theorem 3.10 is somewhat cumbersome. Hence, we decided to first state the essential ingredients and defer the proofs of the helping lemmas to the following subsections. We remark that Lemma 3.12 and Lemma 3.13 hold more general for arbitrary ℓ, whereas $\ell = 3$ is only needed in Lemma 3.11. The following lemma identifies all possible base partitions of $BW_{k,3}$:

Lemma 3.11. *For any odd $k \geq 3$ there are exactly $2 \cdot 2^{\frac{k-1}{2}} + 2^{\frac{k-3}{2}}$ non-isomorphic base partitions of $BW_{k,3}$.*

Next, we extend these base partitions to a proper partition in two steps:

Lemma 3.12. *Let $T_0, \ldots, T_{n-1}$ be a base partition of the edges of $BW_{k,\ell}$. Then there is a unique way to extend it to a partial partition that covers all diagonal edges of distance $d_1, \ldots, d_\ell$.*

Lemma 3.13. *Let $T_0, \ldots, T_{n-1}$ be a partial partition of the edges of $BW_{k,\ell}$ such that exactly all radial edges and all diagonal edges of distance $d_1, \ldots, d_\ell$ are covered, and such that all partial trees are connected, plane, and non-empty. Then, this partial partition can be extended to a partition of $BW_{k,\ell}$ into plane spanning trees. More precisely, there are exactly $2^{\frac{k-1}{2}\ell-1}$ such possible extensions.*

The proofs of these three lemmas can be found in Section 3.3.1, Section 3.3.2, and Section 3.3.3 respectively. We are now ready to prove Theorem 3.10. Let us first show that there cannot be any partition that does not conform to the properties of a base partition:

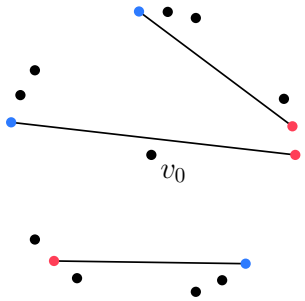

Figure 3.8: Left endpoints are marked blue and right endpoints are red.

Lemma 3.14. *Let $T_0, \dots, T_{n-1}$ be a partition of the edges of $BW_{k,3}$ into plane spanning trees. Then, there are subtrees $T'_0, \dots, T'_{n-1}$ fulfilling the properties of a base partition.*

Proof. Define the subtrees T'_i to consist only of the radial and maximal diagonal edges (of the respective T_i). Properties (a) and (c) are clearly fulfilled. Moreover, every T'_i is non-empty and plane, and connectedness follows from the fact that the maximal diagonal edges define the span and within each span there can only be radial edges (Lemma 3.6). Hence, property (d) holds as well. Finally, the first part of property (b) holds by Proposition 3.9 and by Proposition 3.5 there are no further maximal diagonal edges. □

Theorem 3.10. *For any odd $k \geq 3$, there are exactly $4^{k-1} + 4^{k-2}$ non-isomorphic partitions of $BW_{k,3}$ into plane spanning trees.*

Proof. By Lemma 3.14 it suffices to count the partitions that we get from our construction via base partitions (Lemma 3.11, Lemma 3.12, and Lemma 3.13). These are:

$$\begin{aligned} 2^{3\frac{k-1}{2}-1}(2 \cdot 2^{\frac{k-1}{2}} + 2^{\frac{k-3}{2}}) &= 2^{3\frac{k-1}{2}-1+\frac{k-1}{2}+1} + 2^{3\frac{k-1}{2}-1+\frac{k-3}{2}} \\ &= 2^{2(k-1)} + 2^{2(k-2)} \\ &= 4^{k-1} + 4^{k-2} \end{aligned}$$

This concludes the proof. □

We now present the proofs of Lemma 3.11, Lemma 3.12, and Lemma 3.13, though in reverted order, as we find that more instructive.

For convenience, we associate the endpoints of non-radial edges as left and right viewed from the center v_0 as follows: let $e = uv$ be a non-radial edge such that (v_0, u, v) forms a left-turn, then u is called *right* endpoint of e and v is called *left* endpoint of e; see Figure 3.8.

3.3.1 Proof of Lemma 3.13 (full extension)

Lemma 3.13. *Let $T_0, \dots, T_{n-1}$ be a partial partition of the edges of $BW_{k,\ell}$ such that exactly all radial edges and all diagonal edges of distance $d_1, \dots, d_\ell$ are covered, and such that all partial trees are connected, plane, and non-empty. Then, this partial partition can be extended to a partition of $BW_{k,\ell}$ into plane spanning trees. More precisely, there are exactly $2^{\frac{k-1}{2}\ell-1}$ such possible extensions.*

Proof. First of all note that for any $i \geq \ell$, there are exactly $2n-1$ edges of distance d_i. Starting from $d_i = d_\ell$, we iteratively add all edges of distance d_{i+1} to our partial partition. To this end, consider all edges $e_1, e_2, \dots, e_{2n-1}$ of distance d_i in clockwise circular order. For every e_j there are two edges of distance d_{i+1} in e_j^-, both of which share an endpoint with e_j (one of them with the left endpoint of e_j and the other with the right endpont of e_j). By Lemma 3.7 exactly one of the two edges belongs to the same tree as e_j and initially both choices are valid. However, once the first edge of distance d_{i+1} is assigned to a tree, the distribution of the remaining edges of distance d_{i+1} is determined. More precisely, at the beginning of each iteration, fix an orientation $o \in \{\text{left}, \text{right}\}$. Then, for each edge e_j of distance d_i attach to e_j the edge e_j^o, which is the edge of distance d_{i+1} incident to the endpoint of e_j corresponding to o (see Figure 3.9 for an illustration). Since there are equally many edges of distance d_i and d_{i+1}, every edge of distance d_{i+1} is added to exactly one tree. Furthermore, adding the distance d_{i+1} edges preserves planarity as well as the tree structure. Continue this process until distance $d_{i+1} = d_{\frac{k+1}{2}\ell-1} = 1$, where we have 2 choices (left or right) in each step independently. Then all edges are covered and, since we started with n non-empty partial trees, all T_i's form plane spanning trees.

Lastly, since there are two choices (left/right) in each iteration to pick the next smaller diagonal edge and in total we have $\frac{k+1}{2}\ell - 1 - \ell = \frac{k-1}{2}\ell - 1$ iterations, there are exactly $2^{\frac{k-1}{2}\ell-1}$ possible extensions. □

3.3.2 Proof of Lemma 3.12 (base extensions)

Lemma 3.12. *Let $T_0, \dots, T_{n-1}$ be a base partition of the edges of $BW_{k,\ell}$. Then there is a unique way to extend it to a partial partition that covers all diagonal edges of distance $d_1, \dots, d_\ell$.*

Proof. The proof is very similar to the one of Lemma 3.13, with the difference that we need to go through the distances in each pair of opposite groups separately. So, let $\mathcal{G}, \mathcal{G}'$ be a pair of opposite groups and consider some distance d_i with $i \in \{2, \dots, \ell\}$ (assuming all edges of larger distance between this pair are already covered). Note that by assumption (b) of the base partition, the one edge of distance d_1 is covered (providing a base case). Furthermore, also by assumption (b), for every distance d_i ($i \leq \ell$) there is

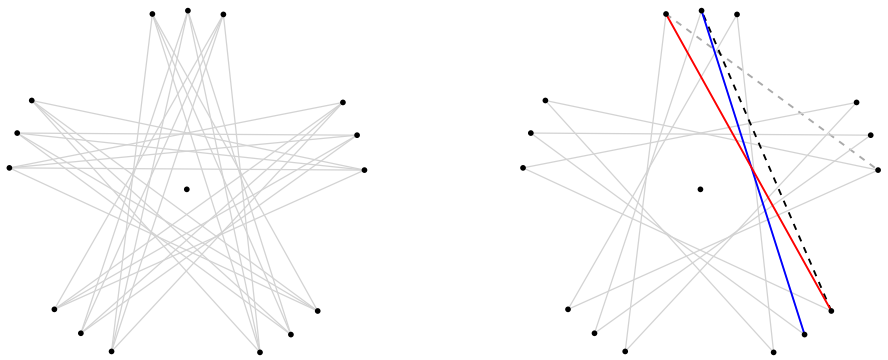

(a) **(b)**

Figure 3.9: (a) All diagonal edges of distance $d_1, \ldots, d_\ell$. (b) All solid edges are of distance d_i and are already covered. If the dashed black edge (of distance d_{i+1}) is attached to the left endpoint of the blue edge, the red edge has only one choice, namely to also attach the next edge (dashed gray) to the left endpoint.

exactly one edge e that is already covered. Assumption (c) implies that e is maximal in its partial tree T_k, in particular, no edge of distance d_{i-1} between $\mathcal{G}$ and $\mathcal{G}'$ belongs to T_k.

Now, at least one of the endpoints of e is incident to an edge f of distance d_{i-1} between $\mathcal{G}$ and $\mathcal{G}'$ (it could be both, then name the other f'). More precisely, let f be incident to the left endpoint of e and f' be incident to the right endpoint of e. Let f be colored blue and f' be colored red. The crucial observation is that e now blocks the left extension of the blue tree as well as the right extension of the red tree. Therefore the red and the blue tree have a unique extension and by using those, they subsequently fix the orientation of the extension for all further trees with an edge of distance d_{i-1} between $\mathcal{G}$ and $\mathcal{G}'$ (to left for all edges to the left of the red edge f' and to right for all edges to the right of the blue edge f – see Figure 3.10 for an illustration). Finally, since there is exactly one more edge of distance d_i than of distance d_{i-1} between $\mathcal{G}$ and $\mathcal{G}'$, this uniquely determines the color for all edges of distance d_i. And similar to the proof of Lemma 3.13, all extended partial trees are still connected, plane, and non-empty. □

3.3.3 Proof of Lemma 3.11 (base partition)

Note that Lemma 3.12 and Lemma 3.13 imply that in order to prove only the existence in Theorem 3.10 it suffices to provide a single base construction for $BW_{k,3}$. We, however, give a full characterization of all possible such base partitions.

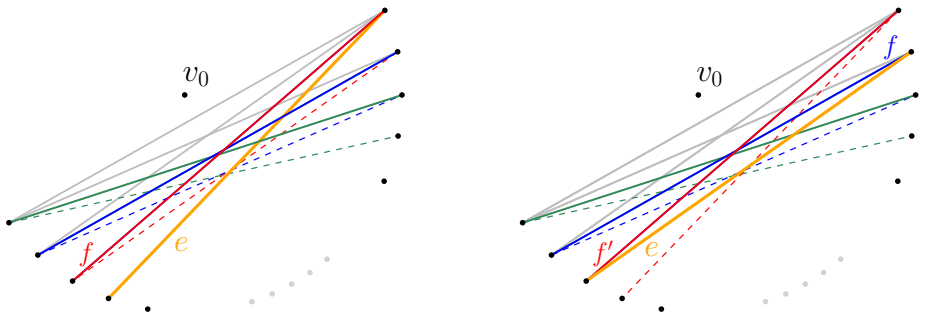

(a) **(b)**

Figure 3.10: Illustration of Lemma 3.12. All edges up to distance $d_{i-1} = d_3$ (solid, non-orange) and one edge e of distance d_i (solid orange) are covered. The edge e now subsequently forces the continuations for the previous trees (dashed). In (a), e is incident to one edge of distance d_{i-1} and in (b) it is incident to two such edges.

Property (b) of a base partition requires us to have one maximal diagonal edge of each distance d_1, d_2, and d_3 between every pair of opposite groups. We need to pair two of them in each tree (except for T_0, that only contains one such edge). The following lemma, which we state in more general terms of subgraphs since we need it that way later (Chapter 4), gives a restriction on which edges we can pair.

Lemma 3.15. *Let $e_1, \dots, e_m$ be pairwise (non-crossing) incomparable edges of a plane subgraph of $BW_{k,\ell}$. Then, $\sum_i \operatorname{dist}(e_i) \leq 2n - 1$ $(= k\ell)$ holds. Moreover, if $m = 2$,*

$$\operatorname{dist}(e_1) + \operatorname{dist}(e_2) \leq 2n - 2$$

holds.

Proof. The first part follows immediately from Observation 3.3 and the second part then from the fact that two edges cannot cover the entire range of the $2n - 1$ boundary edges. □

Lemma 3.16. *For any two incomparable edges with distances d_i and d_j $(1 \leq i, j < \frac{k+1}{2}\ell)$ belonging to the same plane subgraph of $BW_{k,\ell}$, the inequality*

$$i + j \geq \ell + 1$$

holds.

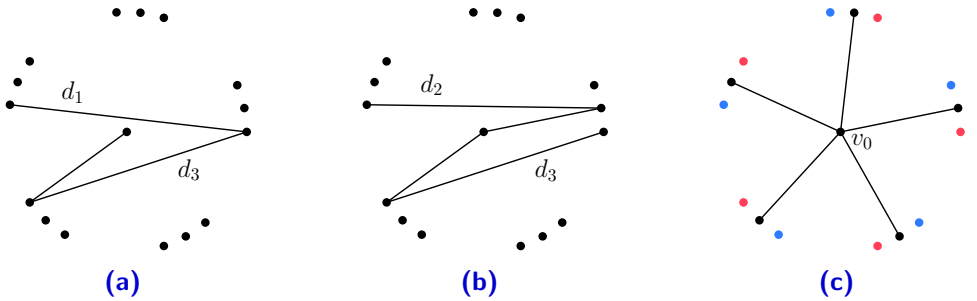

Figure 3.11: (a) Maximal diagonal edges sum to $3k-1$. (b) Maximal diagonal edges sum to $3k-2$. (c) Center radial edges; left vertices are blue and right vertices are red.

Proof. We directly compute:

$$\begin{aligned} i+j &= \frac{k+1}{2}\ell - d_i + \frac{k+1}{2}\ell - d_j \\ &= (k+1)\ell - (d_i + d_j) \\ &\geq k\ell + \ell - (2n-2) && \text{(Lemma 3.15)} \\ &= k\ell + \ell - (k\ell + 1 - 2) \\ &= \ell + 1, \end{aligned}$$

which concludes the proof. □

In the setting of $BW_{k,3}$, Lemma 3.16 implies that we can pair a maximal diagonal edge of distance d_1 only with one of distance d_3 in its plane spanning tree. And a maximal diagonal edge of distance d_2 can only be paired with one of distance d_2 or d_3.

Further note that if the distances of the two maximal diagonal edges sum to $2n-2 = 3k-1$ (i.e., d_1+d_3 or d_2+d_2) the respective spanning tree has exactly one radial edge in its span, and if they only sum to $3k-2$ (d_2+d_3), there are exactly 2 radial edges in the span; see Figure 3.11(a,b).

In the following we call the radial edges incident to the middle vertex of each group *center radial edges*. Also, we call the two outmost vertices of a group the *right* and *left* vertex of that group, respectively, as viewed from v_0 (that is, v_{3i} is the right and $v_{3(i-1)+1}$ the left vertex of group $\mathcal{G}_i$); see Figure 3.11(c) for an illustration.

Finally, we are ready to prove Theorem 3.10, which we restate for easier readability:

Lemma 3.11. *For any odd $k \geq 3$ there are exactly $2 \cdot 2^{\frac{k-1}{2}} + 2^{\frac{k-3}{2}}$ non-isomorphic base partitions of $BW_{k,3}$.*

Proof. Using the structural properties derived on the way to prove Theorem 3.2, we show how to construct all possible base partitions. First of all, we need

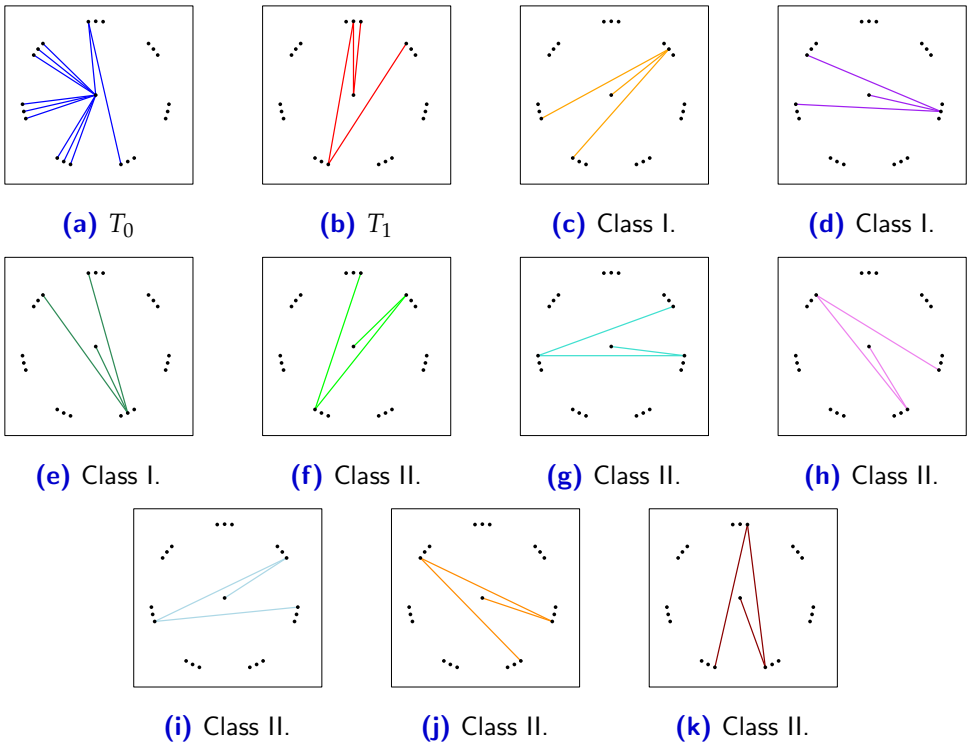

(a) T_0 **(b)** T_1 **(c)** Class I. **(d)** Class I.

(e) Class I. **(f)** Class II. **(g)** Class II. **(h)** Class II.

(i) Class II. **(j)** Class II. **(k)** Class II.

Figure 3.12: Illustration of the different classes of trees and the distribution of maximal diagonal edges. Note that the radial edges are only depicted exemplary and allow certain freedom of choice.

to have a tree T_0 using at least $(3\frac{k-1}{2}+1)$ radial edges and a single maximal diagonal edge e_0 (Lemma 3.4 and Proposition 3.5). Let e_0 again be between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ and let the radial edge $\{v_0, v_{3\frac{k-1}{2}+1}\}$ be part of T_0.

First we argue that $\text{dist}(e_0) = d_1$ holds. If $\text{dist}(e_0) = d_3$, there is one more maximal diagonal edge of distance d_1 than of distance d_3 left, so we cannot pair all edges of distance d_1. Further, if $\text{dist}(e_0) = d_2$, all distance d_1 edges need to be paired with distance d_3 edges and, especially, all remaining distance d_2 edges need to be paired with each other. In particular, the tree containing the center radial edge of $\mathcal{G}_k$ (which cannot be part of T_0) would also have a maximal diagonal edge of distance d_2 between the groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$; a contradiction to Proposition 3.9. Hence, we know $\text{dist}(e_0) = d_1$.

Considering the remaining $k-1$ maximal diagonal edges of distance d_1, they all must be paired with distance d_3 edges. This leaves one distance d_3 edge to be paired with a distance d_2 edge (let the respective tree be T_1) and two distance d_2 edges each for the remaining trees. To summarize (see also Figure 3.12):

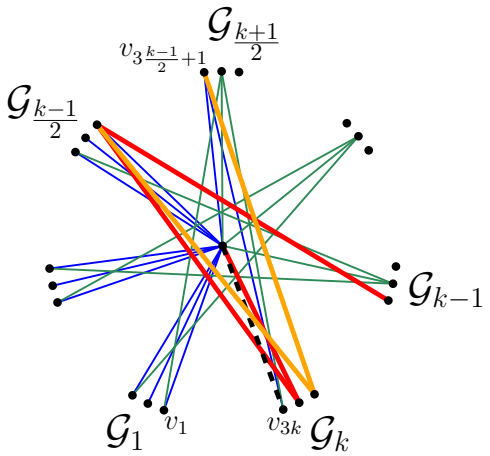

Figure 3.13: Illustration of the claim in the proof of Theorem 3.10. T_0 (blue) uses the d_1 edge between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$. If a Class I tree (green) uses the center radial edge of $\mathcal{G}_{\frac{k+1}{2}}$ (and therefore the d_2 edge between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$), either T_1 (red) or the orange tree uses the d_3 edge between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$. Consequently, the radial edge $\{v_0, v_{3k}\}$ (dashed black) cannot be accommodated anymore.

- T_1 contains one maximal diagonal edge of distance d_2 and one of distance d_3,
- there are $\frac{k-1}{2}$ trees with two distance d_2 maximal diagonal edges (call these trees Class I), and
- $k-1$ trees with one maximal diagonal edge of distance d_1 and d_3 each (call these trees Class II).

Next, consider the $\frac{k+1}{2}$ remaining center radial edges (not used by T_0), which clearly cannot be used by trees of Class II and no tree can use more than one of them. Hence, T_1 and every Class I tree needs to use exactly one such center radial edge. Furthermore, every Class I tree must have the center radial edge incident to its apex and hence, using Proposition 3.9, the center radial edges of the groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ cannot be used by two Class I trees. Therefore, T_1 has to use one of them and some Class I tree T_c the other. The remaining $\frac{k-3}{2}$ Class I trees use the center radial edges of the groups $\mathcal{G}_{\frac{k+3}{2}}$ to $\mathcal{G}_{k-1}$ (with apex in the respective group). Next we show that T_1 uses the center radial edge of $\mathcal{G}_{\frac{k+1}{2}}$ and T_c the one of $\mathcal{G}_k$ (see Figure 3.13 for an illustration):

Claim. *T_c cannot use the center radial edge of $\mathcal{G}_{\frac{k+1}{2}}$.*

Proof. Assume for the sake of contradiction that T_c uses the center radial edge of $\mathcal{G}_{\frac{k+1}{2}}$ and hence, T_1 contains the center radial edge of $\mathcal{G}_k$. Moreover, T_1 has its d_2 edge between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k-1}{2}}$ (since all other d_2 edges are already used by

Class I trees). Then, T_1 has its d_3 edge either between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k+1}{2}}$, or between $\mathcal{G}_{\frac{k-1}{2}}$ and $\mathcal{G}_{k-1}$. In the latter case, the tree using the d_1 edge between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k-1}{2}}$ must have its d_3 edge between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k+1}{2}}$. In any case, the radial edge $\{v_0, v_{3k}\}$ cannot be accommodated anymore, since the corresponding tree would need to have a maximal diagonal edge between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k+1}{2}}$ (all of which are already taken). ■

So, the Class I trees have their apexes in groups $\mathcal{G}_{\frac{k+3}{2}}$ to $\mathcal{G}_k$ and T_1 has to use the last remaining distance d_2 edge (between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_1$). Therefore, the distance d_3 edge of T_1 is either between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ or between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+3}{2}}$.

As a next step, consider the Class II trees. Remember that T_0 uses the distance d_1 edge between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$. Hence, the distance d_3 edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+1}{2}}$ needs to be paired with the distance d_1 edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+3}{2}}$, forcing the respective Class II tree to have its apex in the right vertex of $\mathcal{G}_1$. Subsequently, this forces another Class II tree to have its apex in the right vertex of $\mathcal{G}_2$ and so on, up to group $\mathcal{G}_{\frac{k-1}{2}}$.

Similarly, considering the distance d_3 edge between $\mathcal{G}_{\frac{k-1}{2}}$ and $\mathcal{G}_k$, we get a Class II tree with its apex in the left vertex of $\mathcal{G}_{\frac{k-1}{2}}$. And again subsequently another with its apex in the left vertex of $\mathcal{G}_{\frac{k-3}{2}}$ and so on, down to group $\mathcal{G}_2$. Note that, since T_1 might use the distance d_3 edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+3}{2}}$, we cannot yet conclude where the apex of the last Class II tree will be (let that tree be T_{n-1}). So, let us summarize what we know so far:

- T_0 uses all radial edges $\{v_0, v_i\}$ (for $1 \leq i \leq 3\frac{k-1}{2} + 1$) and the maximal diagonal edge $\{v_{3\frac{k-1}{2}+1}, v_{3k}\}$ of distance d_1 (between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$),
- T_1 uses the center radial edge of group $\mathcal{G}_{\frac{k+1}{2}}$ and the maximal diagonal edge $\{v_1, v_{3\frac{k-1}{2}+2}\}$ of distance d_2 (between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+1}{2}}$),
- the center radial edges of groups $\mathcal{G}_{\frac{k+3}{2}}$ to $\mathcal{G}_k$ are used by Class I trees (with apex in the respective group),
- there are $k-2$ Class II trees with apexes in an outmost vertex from the right vertex of group $\mathcal{G}_1$ to the right vertex of group $\mathcal{G}_{\frac{k-1}{2}}$, and
- the last Class II tree T_{n-1} uses the maximal diagonal edge $\{v_1, v_{3\frac{k+1}{2}}\}$ of distance d_1 (between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+1}{2}}$ — it is the only distance d_1 edge left).

It remains to determine the distance d_3 edge and the second radial edge of T_1 as well as the distance d_3 edge (and hence the apex) of T_{n-1}. We will see that there are three base cases for that:

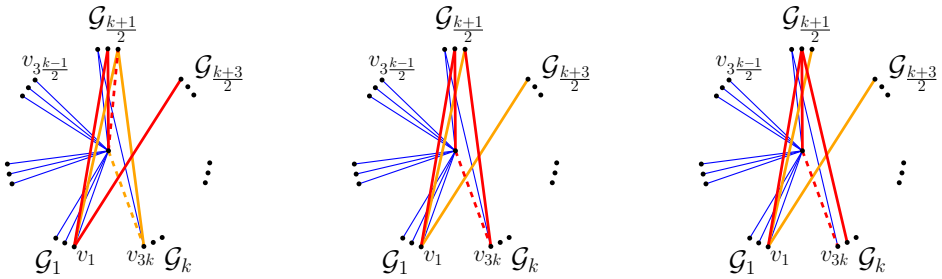

(a) Case 1. **(b)** Case 2a. **(c)** Case 2b.

Figure 3.14: Illustration of the three base cases. T_0 is depicted in blue, T_1 is red, and T_{n-1} orange. The red dashed radial edge depicts the second radial edge, that T_1 has to use.

Case 1: T_1 has its apex in $\mathcal{G}_1$.

That is, T_1 has its distance d_3 edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+3}{2}}$. Hence, the second radial edge of T_1 is $\{v_0, v_{3\frac{k+1}{2}}\}$ and T_{n-1} has its apex in the right vertex of $\mathcal{G}_{\frac{k+1}{2}}$ (see Figure 3.14(a)).

Case 2: T_1 has its apex in $\mathcal{G}_{\frac{k+1}{2}}$.

That is, T_1 has its distance d_3 edge between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$. Hence, T_{n-1} has its distance d_3 edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+3}{2}}$ (so its apex is in the left vertex of $\mathcal{G}_1$) and therefore the second radial edge of T_1 must be $\{v_0, v_{3k}\}$ (since no other tree can use it anymore). However, there are two possibilities for the distance d_3 edge of T_1 now:

a) either $\{v_{3\frac{k+1}{2}}, v_{3k}\}$ as depicted in Figure 3.14(b), or

b) $\{v_{3\frac{k-1}{2}+2}, v_{3(k-1)+2}\}$ as depicted in Figure 3.14(c).

By now, we have fixed all maximal diagonal edges and they fulfill property (b) of a base partition. Finally, it only remains to determine the radial edges of all Class II trees. Note that for each group $\mathcal{G}'$ (from $\mathcal{G}_1$ to $\mathcal{G}_{\frac{k-1}{2}}$) with two apexes of Class II trees in it (at the outmost vertices), those two trees have the same pair of radial edges in their span (in the groups opposite to $\mathcal{G}'$). So we have two (non-isomorphic) choices for each of those cases (independently).

Further, in Case 1 T_{n-1} needs to contain $\{v_0, v_{3k}\}$ and the Class II tree with apex in the right vertex of $\mathcal{G}_1$ needs to contain $\{v_0, v_{3\frac{k+1}{2}+1}\}$ (the other radial edges in its span are used by T_1 and T_0). This leaves a pair of Class II trees for each group from $\mathcal{G}_2$ to $\mathcal{G}_{\frac{k-1}{2}}$. Meanwhile, in Case 2 we have a pair of Class II trees for each group from $\mathcal{G}_1$ to $\mathcal{G}_{\frac{k-1}{2}}$. That is, in Case 2 we have exactly $2^{\frac{k-1}{2}}$ possibilities (for the Class II trees) to conform to the base partition properties,

while in Case 1 we only have $2^{\frac{k-3}{2}}$ possibilities. Given the two choices for T_1 in Case 2, we get

$$2 \cdot 2^{\frac{k-1}{2}} + 2^{\frac{k-3}{2}}$$

base partitions in total. □

In Appendix B, we give some examples of partitions corresponding to the three base constructions.

Chapter 4

Partitions into Plane Subgraphs

In this chapter we generalize our ideas from Chapter 3 to partitions into arbitrary plane subgraphs. We show a negative answer to the following question:

Question 4.1. *Can every complete geometric graph on* $2n$ *vertices be partitioned into* n *plane* subgraphs*?*

Now that we are concerned with partitions into *subgraphs* it would also be viable to allow other values for the number of subgraphs in a partition, or allow complete geometric graphs with an odd number of vertices. However, given the fact that e.g. a point set S in convex position determines a crossing family of size $\lfloor \frac{|S|}{2} \rfloor$, we cannot aim for a smaller size partition in general. On the other hand, allowing more subgraphs, the problem gets trivial once we allow $2n-1$ subgraphs (since any $2n-1$ stars form a proper partition). Hence, whenever we speak of a partition into subgraphs, we refer to a partition into (at most) n subgraphs, thereby also staying close to the notion of the previous chapters.

In Theorem 2.1 we heavily exploited the structure enforced by spanning trees. This is not possible anymore: we cannot make any assumptions on the number of edges, not even about connectedness. The only property we can (and will) exploit is the fact that we still have maximal diagonal edges and radial edges may only be contained in their span (cf. Lemma 3.6).

Recall that the problem of partitioning a geometric graph into plane subgraphs is equivalent to a classic edge coloring problem, where each edge should be assigned a color in such a way that no two edges of the same color cross (of course using as few colors as possible). This problem belongs to a broader question concerning the chromatic number of intersection graphs of geometric objects, a problem that received considerable attention; see e.g. [PKK+14, Dav21, KKN04]. The geometric objects in our setting are line segments, a natural setting, which was also the topic of the CG:SHOP challenge 2022 [CGC].

Contribution. We show that there exist bumpy wheels that cannot even be partitioned into plane subgraphs. In fact, it turns out that allowing arbitrary plane subgraphs instead of plane spanning trees does not help much: the only bumpy wheel that can be partitioned into plane subgraphs but not into plane spanning trees is $BW_{3,5}$.

Theorem 2.2. *For odd parameters $k, \ell \geq 3$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane subgraphs if and only if $\ell > 5$ or ($\ell = 5$ and $k > 3$).*

The remainder of this chapter is dedicated to the proof of Theorem 2.2, where we first focus on the case $\ell > 5$.

Theorem 4.2. *For any odd parameters $k \geq 3$ and $\ell > 5$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane subgraphs.*

The proof is more technical than for spanning trees. Again, we start with some structural results. Recall that the distances of two incomparable edges sum to at most $2n - 2$ (Lemma 3.15). Also recall the notation $d_i = \frac{k+1}{2}\ell - i$ from Section 3.2. The following result is the analogue of Proposition 3.9:

Proposition 4.3. *Let $D_0, \ldots, D_{n-1}$ be a partition of $BW_{k,\ell}$ into $n = \frac{k\ell+1}{2}$ plane subgraphs (if it exists). Then between each pair of opposite groups and for each $1 \leq i \leq \ell$ there are at least i diagonal edges of distance* at least *d_i that are maximal in their subgraph.*

Proof. First note that for every diagonal edge e it holds $\text{dist}(e) \leq d_1$. Furthermore, every diagonal edge e with $\text{dist}(e) \geq d_\ell$ connects opposite groups.

Observe that between every pair of opposite groups and every $1 \leq i \leq \ell$ there is a crossing family of size i (where each edge has distance d_i). Hence, all these edges need to belong to different subgraphs. Furthermore, each such edge gives rise to a maximal diagonal edge in its subgraph: either it is maximal itself or there is a larger edge that is maximal. However, since $i \leq \ell$, larger edges are necessarily between the same pair of opposite groups. Hence, we conclude that there are at least i maximal diagonal edges of distance at least d_i between each pair of opposite groups. □

Rephrasing Proposition 4.3, we may also say that between each pair of opposite groups there is one maximal diagonal edge of distance (at least) d_1, *another* of distance at least d_2, yet *another* of distance at least d_3, and so on until distance d_ℓ.

And since there are exactly k pairs of opposite groups, Proposition 4.3 guarantees at least $k \cdot \ell$ maximal diagonal edges of distance at least d_ℓ in any partition of $BW_{k,\ell}$. For each pair of opposite groups, we distinctly pick for each $1 \leq i \leq \ell$ one of those edges that has distance at least d_i to get precisely $k \cdot \ell$ edges in total, which we call *forced diagonal edges* in the following (or forced edges for short).

These forced edges will take the role of the maximal diagonal edges from the previous chapter. Let E_{forced} be the set of forced diagonal edges, then from the definition it follows:

$$\begin{aligned}\sum_{e \in E_{\text{forced}}} \text{dist}(e) &\geq k \cdot \sum_{i=1}^{\ell} d_i \\ &= k \cdot \sum_{i=1}^{\ell} \left(\frac{k+1}{2}\ell - i \right) \\ &= k \cdot \left(\frac{\ell^2(k+1)}{2} - \frac{\ell(\ell+1)}{2} \right) \\ &= k \cdot (k\ell - 1)\frac{\ell}{2} \end{aligned} \tag{4.1}$$

Note in the following that Lemmas 3.15 and 3.16 hold especially for forced diagonal edges that are contained in the same plane subgraph, since they are maximal and therefore incomparable.

The following proposition is the analogue of Proposition 3.5 generalized to subgraphs.

Proposition 4.4. *Let $D_0, \ldots, D_{n-1}$ be a partition of $BW_{k,\ell}$ into $n = \frac{k\ell+1}{2}$ plane subgraphs (if it exists). Then one subgraph, say D_0, contains exactly one forced diagonal edge and all other $n-1$ subgraphs contain exactly two forced diagonal edges.*

Proof. In total there are $k \cdot \ell$ forced edges and we have $n = \frac{k\ell+1}{2}$ subgraphs. Hence, to prove the statement, we only need to show that no subgraph can contain more than two forced edges. To this end, we consider the two cases $k = 3$ and $k > 3$ separately:

Case 1: $k > 3$.

By Lemma 3.15, in each subgraph the sum of distances of its maximal diagonal edges may not exceed $k\ell$. If there is a subgraph with more than two forced edges (each of distance at least $d_\ell = \frac{k-1}{2}\ell$), the sum of their distances is at least $3 \cdot \frac{k-1}{2}\ell > k\ell$ for $k > 3$.

For the case $k = 3$, we need a more careful analysis.

Case 2: $k = 3$.

Again by Lemma 3.15, if there is a subgraph D' containing three forced edges, then all three need to have distance exactly $d_\ell = \ell$ (for a combined distance of at most 3ℓ) and the subgraph cannot contain any further forced edges. Consider now the three forced edges of distance d_1 (between the three pairs of opposite groups). The subgraphs containing such a forced edge of distance d_1 cannot contain any other forced edge,

which can be seen as follows: By Lemma 3.16, the d_1 edges can only be paired with an edge of distance d_ℓ (which are already used by D'). Hence, we get three subgraphs containing only one forced edge (of distance d_1). Together with D' this leaves $3\ell - 6$ forced edges of distance at least $d_{\ell-1}$ to be covered by the remaining subgraphs (note that by Lemma 3.15 any of the other subgraphs contains at most two of them). Hence, in total we would already need at least

$$4 + \frac{3\ell - 6}{2} = \frac{3\ell + 1}{2} + \frac{1}{2} = n + \frac{1}{2}$$

plane subgraphs to cover all forced edges.

Therefore, we need $n - 1$ subgraphs containing exactly two forced edges and one subgraph (say D_0) containing one forced edge, to cover all $k \cdot \ell = 2n - 1$ forced edges. □

The following proposition is also a generalization of an idea, that we already used towards the proof of Theorem 3.10.

Proposition 4.5. *Let $D_0, \dots, D_{n-1}$ be a partition of $BW_{k,\ell}$ into $n = \frac{k\ell+1}{2}$ plane subgraphs (if it exists) and let e_0 be the forced diagonal edge in D_0 with $\operatorname{dist}(e_0) = d_1 - x_0$ (for some integer $x_0 \geq 0$). Moreover, for each $1 \leq i \leq n-1$, let $x_i \geq 0$ be such that the distances of the two forced diagonal edges e_i and e_i' in D_i sum to $2n - 2 - x_i$. Then,*

$$\sum_{i=0}^{n-1} x_i \leq \frac{\ell - 1}{2}$$

holds.

Proof. By Equation (4.1), we know that

$$\operatorname{dist}(e_0) + \sum_{i=1}^{n-1} \left(\operatorname{dist}(e_i) + \operatorname{dist}(e_i')\right) \geq k \cdot (k\ell - 1)\frac{\ell}{2}$$

holds for the sum of all forced diagonal edges. Plugging in $\operatorname{dist}(e_0) = \frac{k+1}{2}\ell - 1 - x_0$ and $\operatorname{dist}(e_i) + \operatorname{dist}(e_i') = k\ell - 1 - x_i$, yields

$$\frac{k+1}{2}\ell - 1 + (n-1)(k\ell - 1) - \sum_{i=0}^{n-1} x_i \geq k\ell\frac{k\ell - 1}{2}.$$

Finally, plugging in $n = \frac{k\ell+1}{2}$ and rearranging terms

$$\begin{aligned}\sum_{i=0}^{n-1} x_i &\leq \frac{k+1}{2}\ell - 1 + \frac{k\ell-1}{2}(k\ell-1) - k\ell\frac{k\ell-1}{2} \\ &= \frac{k+1}{2}\ell - 1 - \frac{k\ell-1}{2} \\ &= \frac{k\ell+\ell-2-k\ell+1}{2} \\ &= \frac{\ell-1}{2}\end{aligned}$$

provides the desired result. □

Note that, in Proposition 4.5, equality holds if and only if all forced diagonal edges attain there minimal possible distance, i.e., between every pair of opposite groups there is exactly one forced edge of distance *exactly* d_1, exactly one forced edge of distance *exactly* d_2, exactly one forced edge of distance *exactly* d_3, etc. (cf. Proposition 4.3, which was giving only a lower bound). Indeed, the only step in the proof of Proposition 4.5 that used an inequality was the usage of Equation (4.1), which gives equality precisely under the mentioned condition.

Let the forced diagonal edge e_0 in D_0 from now on connect groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ (similar to the maximal diagonal edge of T_0 in Section 3.2). Further note that two endpoints of the two forced diagonal edges e_i and e_i' in D_i ($i \neq 0$) must be contained in a common group $\mathcal{G}$, since otherwise e_i and e_i' would cross. We call the set of vertices in $\text{span}(e_i, e_i')$ that lie in this common group $\mathcal{G}$ the *apex* of D_i (again as in Section 3.2). Especially note that all radial edges in D_i ($i \neq 0$) are contained in $\text{span}(e_i, e_i')$ (cf. Lemma 3.6).

As in Section 3.2, we are using a counting argument to show that certain edges cannot be covered. To this end, it is convenient to introduce the notion of *additional vertices*. Intuitively, consider the span of a subgraph (say with two forced diagonal edges), then this span contains at least three vertices – two endpoints of the forced edges and at least one vertex in the apex (which may be a common endpoint of the forced edges). All other vertices in the span contribute to the additional vertices (except v_0 which is not relevant here). More formally, note that $\text{cl}(e_0^+)$ contains exactly $(\frac{k-1}{2}\ell + 2 + x_0)$ vertices (those of the first $\frac{k-1}{2}$ groups, two outmost vertices, and x_0 extra vertices). Moreover, for $i \neq 0$, $\text{span}(e_i, e_i')$ contains exactly $(2 + 1 + x_i)$ vertices (two outmost vertices, one vertex in the apex, and x_i extra vertices (possibly also in the apex, making its size larger than 1)). We call the set of all extra vertices *additional vertices*; see Figure 4.1(a).

Lemma 4.6. *In any partition of $BW_{k,\ell}$ into $n = \frac{k\ell+1}{2}$ plane subgraphs any inside radial edge of the last $\frac{k+1}{2}$ groups $\mathcal{G}_{\frac{k+1}{2}}, \ldots, \mathcal{G}_k$ is either incident to the apex or to an additional vertex of its subgraph.*

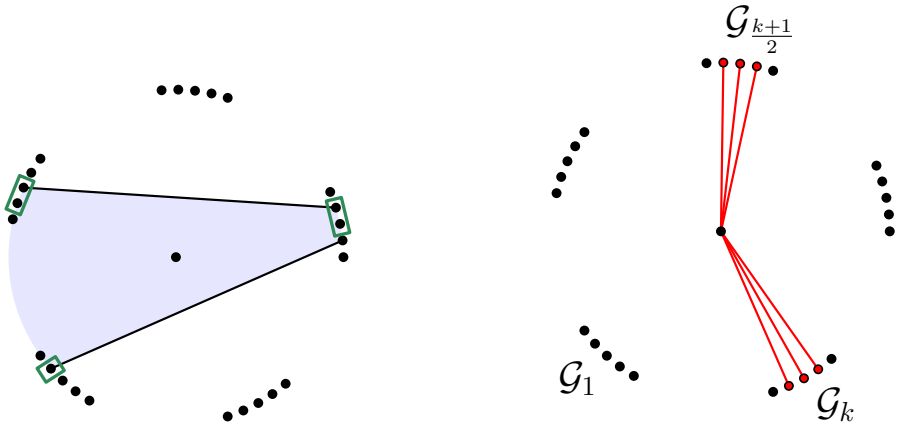

(a) Additional vertices. **(b)** Special radial edges.

Figure 4.1: Illustration of some terms.

Proof. Let f be an inside radial edge of one of the last $\frac{k+1}{2}$ groups. If f belongs to D_0, it is incident to an additional vertex (since it is not in the first $\frac{k-1}{2}$ groups and not incident to one of the two outmost vertices). On the other hand, if f belongs to some D_i with $i \neq 0$, it still cannot be incident to one of the two outmost vertices in $\text{span}(e_i, e_i')$, so it is either incident to the apex or an additional vertex. □

Finally we call the inside radial edges and inside vertices of the two groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ *special radial edges* and *special vertices*; see Figure 4.1(b). Now we are ready to prove Theorem 4.2, which we again restate for easier readability:

Theorem 4.2. *For any odd parameters $k \geq 3$ and $\ell > 5$, the edges of $BW_{k,\ell}$ cannot be partitioned into $n = \frac{k\ell+1}{2}$ plane subgraphs.*

Proof. Note that there are exactly $2\ell - 4$ special radial edges. By Lemma 4.6 they are either incident to an apex (in one of the two groups) or an additional vertex. Any subgraph with an apex in $\mathcal{G}_{\frac{k+1}{2}}$ or $\mathcal{G}_k$ must contain a forced edge between those two groups. By the definition of forced edges there are exactly ℓ of them between $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ and one of them is taken by D_0. Hence, $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ together contain the apexes of at most $\ell - 1$ subgraphs. Finally, Proposition 4.5 gives an upper bound on the total number of additional vertices. Hence, in total we can cover at most

$$(\ell - 1) + \frac{\ell - 1}{2} = \frac{3}{2}(\ell - 1)$$

special radial edges with our n subgraphs (note that we count only one vertex for each apex because if an apex is larger than 1, that is, contains additional

vertices, we already accounted for that in the additional vertices bound). Using $\ell > 5$, we conclude that not all $2\ell - 4$ special radial edges can be covered:

$$\frac{3}{2}(\ell - 1) < \frac{3}{2}(\ell - 1) + \frac{1}{2}(\ell - 5) = 2\ell - 4.$$

□

For the case $\ell = 5$, we need to analyze the structure of our plane subgraphs a little further.

Theorem 4.7. *For any odd parameter $k \geq 5$, the edges of $BW_{k,5}$ cannot be partitioned into $n = \frac{5k+1}{2}$ plane subgraphs.*

Proof. As before, we first consider the special radial edges, that is, the inside radial edges of the groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$. Since $\ell = 5$, there are 6 of them here. Furthermore, any subgraph that has its apex in $\mathcal{G}_{\frac{k+1}{2}}$ or $\mathcal{G}_k$ must use a forced edge between those two groups (the blue stripe in Figure 4.2). However, because D_0 already uses one of these 5 forced edges, there are at most 4 apexes in $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ together. Moreover, Proposition 4.5 yields at most $\frac{5-1}{2} = 2$ additional vertices in total.

To summarize, between the two groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$ there are 6 special radial edges to be covered in 4 apexes. This implies that two of these special vertices must be additional vertices. In particular, there cannot be any further additional vertex in another apex. Moreover, all forced diagonal edges attain their minimal possible distance (see the remark after Proposition 4.5), i.e., we have exactly one forced diagonal edge of each distance $d_1, \ldots, d_5$ between every pair of opposite groups.

Consider now the inside radial edges in all groups from $\mathcal{G}_{\frac{k+3}{2}}$ to $\mathcal{G}_{k-1}$. By Lemma 4.6 and the fact that all additional vertices are special vertices, they must be incident to some apex. Also, since the opposite groups are between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k-1}{2}}$, it is not possible to place special vertices as additional vertices in the respective subgraphs. Hence, these subgraphs use up all forced edges of distances d_2, d_3, d_4 (in the grey stripes in Figure 4.2), except of course those between the pairs of opposite groups $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+1}{2}}$, $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$, and $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k-1}{2}}$.

Since $k \geq 5$, each edge between $\mathcal{G}_1$ and $\mathcal{G}_{\frac{k+1}{2}}$ crosses every edge between $\mathcal{G}_k$ and $\mathcal{G}_{\frac{k-1}{2}}$ (the red stripes in Figure 4.2). Furthermore, two forced edges between the same pair of opposite groups cannot be in the same subgraph (because they are both maximal). Hence, we have 6 forced edges of distances d_2, d_3, d_4 (call them *leftover edges*) between those two pairs of opposite groups, that we still need to pair with a second forced edge in their subgraph.

However, by Lemma 3.16 all forced edges of distance d_1 (except one if used by D_0) need to be paired with a distance d_5 forced edge in their subgraph.

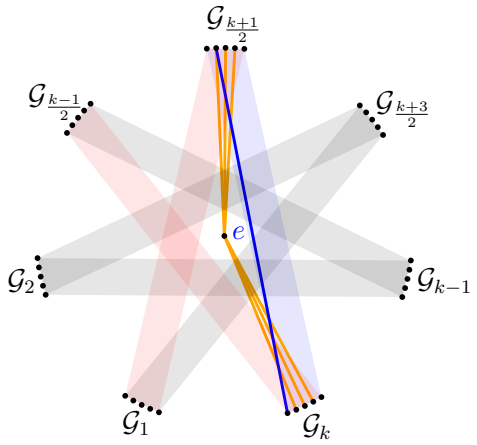

Figure 4.2: High level overview of the proof of Theorem 4.7. We have at most $\frac{5-1}{2} = 2$ additional vertices in total and the blue stripe (which contains the single forced edge e of D_0) has to use both of them. Then, in the grey stripes we must use all forced edges of distances d_2, d_3, d_4. However, since the two red stripes intersect ($k \geq 5$), there will not be enough forced edges to pair all 6 forced edges of distances d_2, d_3, d_4 from the red stripes.

This leaves the 3 forced edges of distances d_2, d_3, d_4 between groups $\mathcal{G}_{\frac{k+1}{2}}$ and $\mathcal{G}_k$, and possibly one forced edge of distance d_5 to pair the leftover edges with. That is two less than what we would need. □

Theorems 4.2 and 4.7 prove the "if" direction of Theorem 2.2. Towards the other direction, using Theorem 2.1, it only remains to show that there is a partition for $BW_{3,5}$. However, in Section 3.1 (see Figure 3.2) we already gave such a partition.

Chapter 5

Generalized Wheels

In this chapter we generalize our construction to non-regular wheel sets. We give a necessary condition in the setting of plane spanning trees (Theorem 2.3) and a full characterization for partitioning into plane double stars (Theorem 2.4). Recall that, for a tuple $N = [n_1, \ldots, n_k]$ of integers $n_i \geq 1$, GW_N denotes the generalized wheel with group sizes n_i. In Section 5.1, we prove the following result:

Theorem 2.3. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning trees if each family of $\frac{k-1}{2}$ consecutive groups contains (strictly) less than $n-2$ vertices.*

Note that the geometric regularity of generalized wheels is not strictly required (but eases the proofs). In fact, we show that for every wheel graph there exists a generalized wheel graph with the same rotation system.

Considering the other side of the story, we show that many generalized wheels can already be partitioned into plane double stars. In fact, we give the following characterization:

Theorem 2.4. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning double stars if and only if there are three families of $\frac{k-1}{2}$ consecutive groups, such that (i) every family contains at most $n-2$ vertices and (ii) every group is in at least one family.*

We phrased Theorem 2.4 this way to make it consistent with Theorem 2.3; however, let us rephrase it in a way that better indicates the gap between the two theorems. Let F_i denote the family of $\frac{k-1}{2}$ consecutive groups starting at $\mathcal{G}_i$ in clockwise order (whenever speaking of a family without further specification, we refer to such a family of $\frac{k-1}{2}$ groups for the remainder of this section). Two families F_i and F_{i+1} are called *consecutive* and $|F_i|$ denotes the number of vertices in F_i. If $|F_i| \leq n-2$ holds, we call F_i *small*, and otherwise *large*.

Corollary 5.1. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N can be partitioned into plane spanning double stars if and only if there are $\frac{k-1}{2}$ consecutive families each containing (strictly) more than $n-2$ vertices.*

Proof. If, for the one direction, there are $\frac{k-1}{2}$ large consecutive families, then there is a group $\mathcal{G}^\star$ (namely the one that is contained in all these $\frac{k-1}{2}$ families) such that any family containing $\mathcal{G}^\star$ is large. In particular, there cannot be three small families covering all groups. Hence, by Theorem 2.4, there is a partition into plane double stars.

On the other hand, if there are no $\frac{k-1}{2}$ large consecutive families, we can find three small families as follows. Note first that every group is contained in some small family. Pick a small family F arbitrarily and let $\mathcal{G}$ be the first group after F (in clockwise order). Among all small families containing $\mathcal{G}$, pick the one that is "furthest" from F, that is, has least overlap with F, and call it F'. Let $\mathcal{G}'$ again be the first group after F' and among all small families containing $\mathcal{G}'$ pick the one furthest from F' and call it F''. Since F'' cannot contain $\mathcal{G}$, we conclude that the three small families F, F', F'' cover all groups. □

5.1 Plane Spanning Trees

Theorem 2.3 extends Theorem 3.2 to generalized wheels and, as we will see, to a large extent the proof is analogous. We remark that the condition in Theorem 2.3 is only a sufficient condition but not a necessary one. In fact, we found some generalized wheels not fulfilling the condition, that still cannot be partitioned into plane spanning trees (verified by computer assistance), for example, $GW_{[2,3,3,4,5]}$ cannot be partitioned.

In Section 3.2 we considered the more restrictive setting of bumpy wheels and showed for which parameters these cannot be partitioned into plane spanning trees. We found this presentation more instructive, however, now we need some of the technical results in the context of generalized wheels. All results from Observation 3.3 to Proposition 3.9 hold for generalized wheels as well, exactly the way they were stated for bumpy wheels, *with two exceptions*: Lemma 3.4 (iii) and Proposition 3.9. In the following, we explain precisely which parts of the mentioned results do not carry over and how to resolve this (also in the proofs of the other results that use Lemma 3.4 (iii) or Proposition 3.9).

Adjusting Lemma 3.4 (iii). Recall that Lemma 3.4 (iii) guarantees, for a spanning tree T_0 with exactly one maximal diagonal edge, the presence of all radial edges along $\frac{k-1}{2}$ consecutive groups plus one. This statement carries one-to-one over to generalized wheels. However, the precise bound $(\frac{k-1}{2}\ell + 1)$ on the number of radial edges does not carry over. We used this value in the proof of Proposition 3.5 to show that there is a single spanning tree in the partition with exactly one maximal diagonal edge.

We claim, however, that this statement also holds in our setting of generalized wheels, in other words, Proposition 3.5 persists. The argument in the proof of Proposition 3.5 can be adjusted as follows. Assume that there are two trees T_0, T_1 with only a single maximal diagonal edge, then T_0 and T_1 together cover all radial edges of $k-1$ groups plus two (using the first part Lemma 3.4 (iii)). Hence, there is only a single group, say $\mathcal{G}_k$, whose radial edges are not entirely covered by T_0 and T_1. Since each of the remaining $n-2$ trees contains at least one radial edge, it follows that $\mathcal{G}_k$ contains at least $n-2+2=n$ vertices.

However, in this case, the condition of Theorem 2.3 is not fulfilled. Or, phrased the other way around, a generalized wheel that fulfills the condition of Theorem 2.3, also fulfills Proposition 3.5.

Adjusting Proposition 3.9. It requires some effort to fit the formulation of Proposition 3.9 to generalized wheels. In essence, this reformulation only affects the *values* of the distances of maximal edges (which is not $\left(\frac{k+1}{2}\ell - j\right)$ anymore), however, not the value of the distance is relevant but the number of different distances. We give the precise formulation of Proposition 3.9 in the context of generalized wheels and add the proof, which is very similar to the one of Proposition 3.9, for the sake of completeness:

Proposition 5.2 (cf. Proposition 3.9)**.** *Let $T_0, \dots, T_{n-1}$ be a partition of GW_N (with $N = [n_1, \dots, n_k]$ positive integers) into plane spanning trees (if it exists). Further let $\mathcal{G}_x, \mathcal{G}_y$ be a pair of opposite groups and $d^\star$ be the distance of a largest edge between $\mathcal{G}_x$ and $\mathcal{G}_y$. Then, for each $j \in \{1, \dots, \max(n_x, n_y)\}$ there is at most one diagonal edge connecting $\mathcal{G}_x$ and $\mathcal{G}_y$ of distance $d'_j = (d^\star + 1 - j)$ that is maximal in its tree.*

Proof. Without loss of generality, let $n_x \geq n_y$, i.e. $\max(n_x, n_y) = n_x$. Observe that for any $j \in \{1, \dots, n_x\}$ there are *at most* j edges of distance d'_j between $\mathcal{G}_x$ and $\mathcal{G}_y$ (since some edges of this distance may be between a different pair of groups now); see Figure 5.1(a).

For the remainder of the proof, we only consider edges with one endpoint in $\mathcal{G}_x$ and the other endpoint in $\mathcal{G}_y$. First note that for every $j \in \{1, \dots, n_x\}$, there are $j^\star = \min(j, n_y)$ edges of distance d'_j (between $\mathcal{G}_x$ and $\mathcal{G}_y$).

Consider for some $j \in \{2, \dots, n_x\}$ the distance d'_j and let $c_1, \dots, c_{j^\star}$ be the colors used for the $j^\star$ edges of this distance, which are all distinct, since these edges form a crossing family.

Case 1: $j \leq n_y$, i.e., $j^\star = j$.

By Lemma 3.7, the $j-1$ pairwise crossing edges of (the larger) distance d'_{j-1} must use the same color as an edge of distance d'_j (w.l.o.g. $c_1, \dots, c_{j-1}$). Hence, the corresponding edges of distance d'_j cannot be maximal (again using the argument that there cannot be two maximal

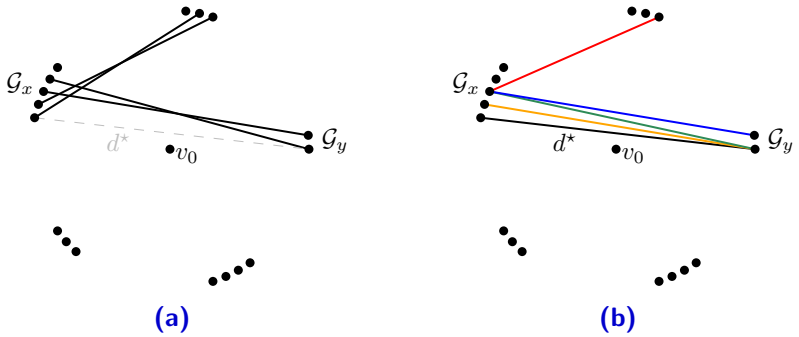

Figure 5.1: Illustration of the proof of Proposition 5.2. (a) The edges of distance d'_j ($j = 4$ here) form a crossing family. (b) Maximal diagonal edges may not all be between $\mathcal{G}_x$ and $\mathcal{G}_y$.

diagonal edges between the same pair of groups). On the other hand, the color c_j cannot be used by any edge of larger distance, since again by Lemma 3.7 there would also be an edge of color c_j with distance d'_{j-1}. Hence, indeed the only edge of distance d'_j that is maximal in its tree is the one of color c_j.

Case 2: $j > n_y$, i.e., $j^\star = n_y$.

First note that now there are equally many edges of distance d'_j and d'_{j-1}, namely n_y many. Using the arguments as in Case 1, we conclude that there is no edge of distance d'_j that is maximal in its tree.

Lastly, for $j = 1$ observe that the single edge of distance d'_1 is clearly maximal. □

Putting everything together. We are now ready to prove Theorem 2.3. For convenience we introduce a little more notation and slightly rephrase the theorem. For $j \in \{1, \dots, k\}$, define (the indices of) the $\frac{k-1}{2}$ consecutive groups starting at $\mathcal{G}_j$ as I_j, that is, $I_j = \{j, j+1, \dots, j+\frac{k-1}{2}-1\}$ (as usual all indices are taken modulo k). Then, Theorem 2.3 can be equivalently phrased as follows: GW_N cannot be partitioned into plane spanning trees if for all $j \in \{1, \dots, k\}$ the inequality $\sum_{i \in I_j} n_i < n - 2$ holds.

Proof of Theorem 2.3. First of all note that if there is an i such that $n_i = 1$ (that is, a group consisting of only one vertex v_1), the condition of Theorem 2.3 cannot be satisfied: Consider the line through v_0 and v_1, then on each side there are $\frac{k-1}{2}$ consecutive groups and one side must contain a total of at least $n-1$ vertices. So, we can assume from now on that $n_i \geq 2$ holds for all $1 \leq i \leq k$. The remainder of the proof is analogous to the proof of Theorem 3.2.

Let the condition in Theorem 2.3 be satisfied and assume for the contrary that there was a partition $T_0, \dots, T_{n-1}$. Again, since T_0 uses at least k outmost radial edges (Proposition 3.5), we can have at most $k+1$ spanning trees containing an outmost radial edge.

Counting the trees not containing any outmost radial edge, we need to be a bit more careful now. We claim that there are at most

$$\max_{j\in\{1\dots k\}} \sum_{i\in I_j}(n_i - 2)$$

spanning trees not containing any outmost radial edge. The arguments are analogous as in Theorem 3.2: By Proposition 3.8, the apex of a special wedge can only be in one of the $\frac{k+1}{2}$ groups whose radial edges are not fully used by T_0. Also by Proposition 3.8 each such group $\mathcal{G}_i$ contains the apex of at most $n_i - 2$ special wedges. Furthermore, by Proposition 5.2, instead of considering $\frac{k+1}{2}$ consecutive groups it suffices to sum over $\frac{k-1}{2}$ consecutive groups. The crucial difference to Theorem 3.2 is that we cannot just assume w.l.o.g. a position of T_0 and hence, we need to take the maximum over all possible families of $\frac{k-1}{2}$ consecutive groups.

Hence, whenever

$$k + 1 + \sum_{i\in I_j}(n_i - 2) < n$$

holds for all j, we cannot find enough spanning trees. Rearranging terms, this inequality is equivalent to $\sum_{i\in I_j} n_i < n - 2$ (recall that $|I_j| = \frac{k-1}{2}$). □

5.2 Dropping the geometric regularity

In this section, we illustrate how to drop the geometric regularity of generalized wheel sets. More precisely, we show that for every wheel set there exists a weakly isomorphic generalized wheel, i.e., a generalized wheel with the same rotation system.

Similarly to Chapter 3 we will use the notation e^- for the (open) halfplane defined by (the supporting line through) e and not containing v_0 (the only vertex inside the convex hull). Note that we do not need an even number of vertices here, so we consider point sets on n vertices now (instead of $2n$).

Theorem 5.3. *Let P be a set of $n \geq 4$ points in the plane in general position with exactly one point v_0 inside the convex hull. Then there exists a generalized wheel GW_N defining the same rotation system.*

Proof. Denote the vertices on the convex hull by $v_1, \dots, v_{n-1}$ in clockwise order (around v_0) and for each vertex v_i ($1 \leq i \leq n-1$) define the *opposite boundary edge* (denoted by $\bar{e}(v_i)$) as the unique boundary edge $\{v_j, v_{j+1}\}$ such that $v_0 \in \Delta v_i v_j v_{j+1}$. Further define the groups as sets of vertices having

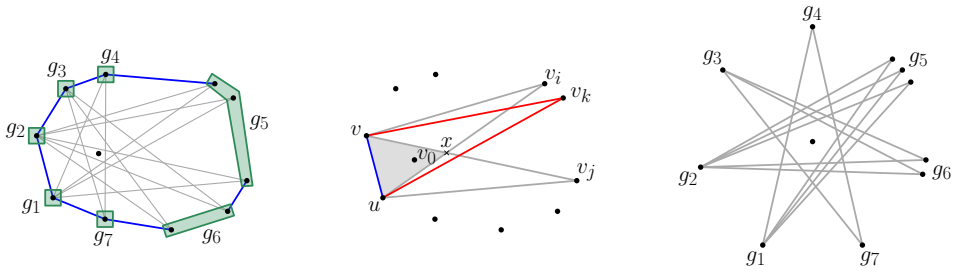

Figure 5.2: (a) Defining the groups (marked in green); opposite boundary edges drawn blue. (b) For any vertex v_k between v_i and v_j, the triangle Δxuv has to be contained in $\Delta v_k uv$, since v_k lies in the cone $v_i x v_j$. (c) The corresponding generalized wheel.

the same opposite boundary edge (see Figure 5.2(a)). We will see that these groups correspond precisely to the groups in the generalized wheel. Note that not all vertices can be in the same group (for example, the endpoints of $\bar{e}(v_i)$ cannot be in the same group as v_i).

First, we show that each of those groups consists of consecutive vertices in P (along the convex hull). Let v_i, v_j $(i < j)$ be two vertices on the convex hull with $\bar{e}(v_i) = \bar{e}(v_j) := \{u, v\}$. We claim that all vertices of P in $e^- = \{v_i, v_j\}^-$ (w.l.o.g. $v_{i+1}, \ldots, v_{j-1}$) belong to the same group as v_i and v_j. Indeed, for any v_k with $i < k < j$ we have $\Delta xuv \subseteq \Delta v_k uv$ (where x is the intersection of $\{v_i, u\}$ and $\{v_j, v\}$); see Figure 5.2(b). And since $v_0 \in \Delta xuv$, the claim follows.

Next, we show that for each opposite boundary edge $\bar{e}(v_k) = \{u, v\}$ (of some hull vertex v_k) the vertices u and v belong to different groups. Indeed, since $v_0 \in \Delta v_k uv$, we get that $\bar{e}(u)$ must lie in $\mathrm{cl}(e_v^-)$ of the edge $e_v = \{v, v_k\}$ (see also Figure 5.2(b)). Similarly, $\bar{e}(v)$ must lie in $\mathrm{cl}(e_u^-)$ of the edge $e_u = \{u, v_k\}$. Hence, u and v have different opposite boundary edges.

Further, we show the following two properties by induction on the number of vertices $m = n - 1$ on the convex hull:

(P1) P defines an odd number of groups and

(P2) for each v_i, its opposite boundary edge $\bar{e}(v_i)$ splits the remaining groups into equal parts with respect to the number of groups, that is, the line through v_i and any point of (the interior of) $\bar{e}(v_i)$ has equally many groups on each side (excluding the group containing v_i).

Concerning the base case $m = 3$ note that each of the 3 vertices forms its own group and hence, (P1) and (P2) hold. So, let P be a point set with $m \geq 4$ hull vertices and consider P' by removing an arbitrary hull vertex v_i such that $P' = P \setminus v_i$ still contains exactly one point inside the convex hull

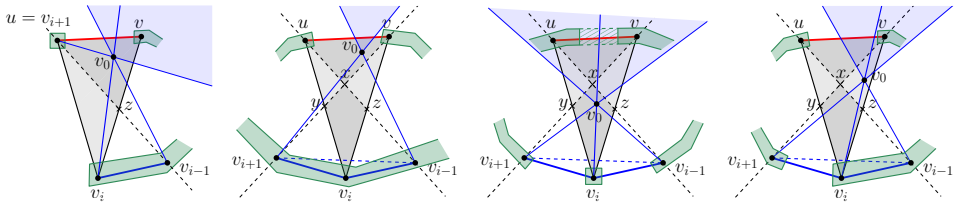

(a) Case 0. **(b)** Case 1. **(c)** Case 2. **(d)** Case 3.

Figure 5.3: Illustration of the case distinction. The blue shaded region always depicts the area of vertices having opposite boundary edge $\{v_{i-1}, v_{i+1}\}$.
(a) Case 0: v_i joins the group of v_{i-1} and all remaining groups remain unchanged.
(b) Case 1: v_i joins the group of v_{i-1} and v_{i+1} and all remaining groups remain unchanged.
(c) Case 2: v_i forms a new group of size 1 and the group containing u and v is split into two groups, one with opposite boundary edge $\{v_{i-1}, v_i\}$ and the other with opposite boundary edge $\{v_i, v_{i+1}\}$.
(d) Case 3: v_i joins the group of v_{i-1} and all remaining groups remain unchanged.

(that is, $v_0 \notin \Delta v_i v_{i+1} v_{i-1}$). By the induction hypothesis, P' fulfills properties (P1) and (P2). Now, insert v_i back in and let $\bar{e}(v_i) = \{u, v\}$. First, we will discuss the case that u and v_{i+1} (or analogously v and v_{i-1}) coincide (Case 0); otherwise, consider the (crossing) edges $\{v_{i-1}, u\}$ and $\{v_{i+1}, v\}$ and call their intersection x; further denote their intersections with $\Delta v_i u v$ by y and z (see Figure 5.3). Then, there are 4 different regions of $\Delta v_i u v$ that may contain v_0 (shaded gray in Figure 5.3(b-d)). We consider these cases separately (note that in Case 0 only one of the two crossing points y and z exists, say z):

Case 0: u and v_{i+1} coincide (or analogously v and v_{i-1} coincide).

First, if u and v_{i+1} coincide, v and v_{i-1} cannot coincide too, since otherwise there would only be three vertices on the convex hull. Furthermore, v_0 is contained in the triangle uzv, since otherwise v_0 would not lie in the interior of the convex hull without v_i.

Then v_i and v_{i-1} have the same opposite boundary edge $\{u, v\}$, i.e., v_i joins the group of v_{i-1}. Moreover, the edge $\{v_{i-1}, v_{i+1}\}$ is replaced by the two edges $\{v_{i-1}, v_i\}$, $\{v_i, v_{i+1}\}$. There cannot be any vertex having $\{v_{i-1}, v_i\}$ as opposite boundary edge and all vertices that had opposite boundary edge $\{v_{i-1}, v_{i+1}\}$ now have opposite boundary edge $\{v_i, v_{i+1}\}$; see Figure 5.3(a). In total, the number of groups remains unchanged, hence (P1) holds. Concerning (P2) note that all groups remain unchanged except v_i joining the group of v_{i-1}. Furthermore, all vertices that had opposite boundary edge $\{v_{i-1}, v_{i+1}\}$ now have opposite boundary edge $\{v_i, v_{i+1}\}$. Hence, using the induction hypothesis, (P2) also holds.

Case 1: $v_0 \in \Delta xuv$.

Here v_{i-1} and v_{i+1} belong to the same group with opposite boundary edge $\{u, v\}$. So, v_i joins this already existing group. Moreover, the edge $\{v_{i-1}, v_{i+1}\}$ is replaced by the two edges $\{v_{i-1}, v_i\}$, $\{v_i, v_{i+1}\}$. However, due to the convex position of the hull vertices there cannot be any vertex having $\{v_{i-1}, v_i\}$ or $\{v_i, v_{i+1}\}$ as opposite boundary edge (see Figure 5.3(b)). In total, the number of groups remains unchanged, hence (P1) holds. Further (P2) holds by induction hypothesis (clearly for any vertex other than v_i, and for v_i since it has the same opposite boundary edge as v_{i-1} and v_{i+1}).

Case 2: $v_0 \in \Diamond v_i yxz$.

In this case, $\bar{e}(v_{i-1})$ lies in e_u^- of the edge $e_u = \{u, v_i\}$ and $\bar{e}(v_{i+1})$ lies in e_v^- of the edge $e_v = \{v, v_i\}$. Hence, v_{i-1} and v_{i+1} are in different groups already in P' and in addition v_i forms a new group of size one in P. Furthermore, u and v belong to the same group in P' (with opposite boundary edge $\{v_{i-1}, v_{i+1}\}$). However, after inserting v_i, this group will be split into two parts — the u part with opposite boundary edge $\{v_i, v_{i-1}\}$ and the v part with opposite boundary edge $\{v_i, v_{i+1}\}$ (see Figure 5.3(c)). All other groups stay the same. So in total, the number of groups increases by two, that is, it remains odd (confirming (P1)).

Concerning (P2), each vertex in the group containing u gains precisely one group on each side of its line through the new opposite boundary edge $\{v_i, v_{i-1}\}$, namely the group containing v_i on the one side and the group containing v on the other side. The same holds for each other vertex in e_u^- (without any change to the opposite boundary edge). Similarly, the vertices in the group containing v and all other vertices in e_v^- gain exactly one additional group on each side as well. Finally, v_i fulfills (P2) because (P2) holds for u (and v) in P' and any line through v_i and (the interior of) $\{u, v\}$ has the same groups on each side in P as any line through u and (the interior of) $\{v_{i-1}, v_{i+1}\}$ in P', with the addition of the group containing u on the one side and the group containing v on the other side. Hence, also (P2) holds for all vertices.

Case 3: $v_0 \in \Delta xvz$.

Similar as before, v_{i-1} and v_{i+1} belong to different groups already in P'. However, v_i now joins the group of v_{i-1} in P. Moreover, the group that had $\{v_{i-1}, v_{i+1}\}$ as opposite boundary edge (that is, the group containing v), now has the opposite boundary edge $\{v_i, v_{i+1}\}$ and there is no vertex having $\{v_{i-1}, v_i\}$ as opposite boundary edge (see Figure 5.3(d)). Therefore, the total number of groups remains unchanged (confirming (P1)). Concerning (P2) note that all groups remain unchanged except v_i joining the group of v_{i-1}. Furthermore, all

vertices that had opposite boundary edge $\{v_{i-1}, v_{i+1}\}$ now have opposite boundary edge $\{v_i, v_{i+1}\}$. Hence, using the induction hypothesis, (P2) also holds.

Case 4: $v_0 \in \Delta xyu$.

This is analogous to Case 3 (just laterally reversed).

Finally, by property (P1) we can define the generalized wheel GW_N (recall that we need an odd number of groups here) with the exact same group sizes in the same circular order as defined for P (see Figure 5.2(c)).

It remains to argue that P and GW_N define the same rotation system, i.e., every point encounters the remaining points in the same cyclic order. In particular we need to argue that v_0 always appears at the "right" spot: By property (P2) and the definition of generalized wheels we know that for each extreme point v_i, the opposite boundary edge $\bar{e}(v_i)$ of v_i is the same for P and GW_N. Furthermore, in the cyclic order around v_i, the inner point v_0 appears between the two endpoints of $\bar{e}(v_i)$. Hence, P and GW_N define the same rotation system, which concludes the proof. □

5.3 Plane Double Stars

This section is dedicated to the proof of Theorem 2.4.

Theorem 2.4. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning double stars if and only if there are three families of $\frac{k-1}{2}$ consecutive groups, such that (i) every family contains at most $n-2$ vertices and (ii) every group is in at least one family.*

Roughly speaking this theorem tells us that generalized wheels with rather evenly distributed group sizes cannot be partitioned into double stars, whereas those with very uneven group sizes can be. An exception are *regular* wheels, where every family of $\frac{k-1}{2}$ consecutive groups contains precisely $n-1$ vertices and hence, can always be partitioned into plane double stars.

Towards the proof of Theorem 2.4 we need a couple of definitions and results introduced by Schnider [Sch16]. The edge v_1v_2 in a double star (where v_1 and v_2 are the non-leaf vertices) is called *spine edge*[1]. Given a partition of the edge set of a complete geometric graph into double stars, the collection of spine edges of the double stars forms a perfect matching [Sch16, Lemma 2], called the *spine matching*. Conversely, a perfect matching M (on a point set P) for which there exists a partition of $K(P)$ into plane double stars such that the edges of M are the spines of the double stars is also called *spine matching*.

[1]The term "spine edge" is a frequently used term and in later chapters we will also define it in other contexts.

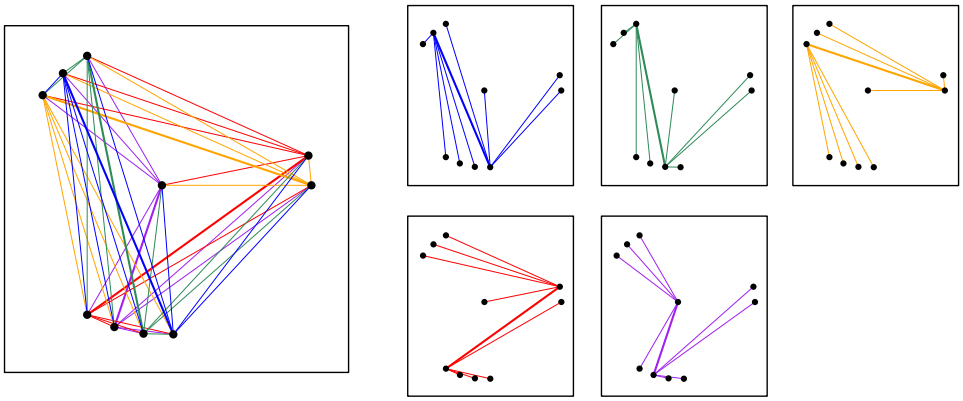

Figure 5.4: A partition of the generalized wheel $GW_{[2,3,4]}$ into plane double stars. The thicker edges form a spine matching.

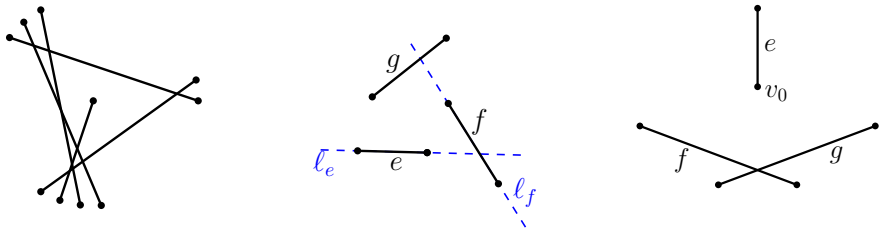

(a) Spine matching (see also Figure 5.4).

(b) Stabbing chain.

(c) Cross-blocker.

Figure 5.5: Illustration of some terms.

Let e and f be two non-adjacent edges and let s be the intersection of their supporting lines. Recall that we say e and f *cross*, if s lies in both e and f. If s lies in f but not in e, we say that e *stabs* f and we call the vertex of e that is closer to s the *stabbing vertex* of e. If s lies neither in e nor in f, or even at infinity, we say that e and f are *parallel*. A *stabbing chain* are three edges, e, f and g, where e stabs f and f stabs g. We call f the middle edge of the stabbing chain. Note that a stabbing chain implies two interior points, so in our setting of wheel sets, there are no stabbing chains.

A *cross-blocker* is a perfect matching on six points in wheel position, where the edge e connecting to the interior point v_0 stabs both other edges f and g, f and g cross, and v_0 is not in the convex hull of f and g. See Figure 5.4 and Figure 5.5 for an illustration of these terms.

Schnider [Sch16, Theorem 9] showed that a spine matching can neither contain two parallel edges nor a cross-blocker.[2] On the other hand, it turns out that for wheel sets, these two configurations are the only two obstructions:

[2]There is a third configuration that cannot occur, but this configuration requires a stabbing chain, so when only considering wheel sets, we may ignore this.

Theorem 5.4 ([Sch16, Theorem 11])**.** *Let P be a point set in general position and let M be a perfect matching on P, such that*

(a) *no two edges are parallel,*

(b) *if an edge e stabs two other edges f and g, then the respective stabbing vertices of e lie inside the convex hull of f and g, and*

(c) *if there is a stabbing chain, then the stabbing vertex of the middle edge lies inside the convex hull of the other two edges.*

Then M is a spine matching.

In the setting of wheel sets, case (c) cannot occur, whereas cases (a) and (b) correspond exactly to the obstructions mentioned above. We thus get the following characterization of generalized wheel sets that allow a partition into double stars as a corollary of Theorem 2.4:

Corollary 5.5. *A generalized wheel GW_N can be partitioned into plane spanning double stars if and only if it admits a perfect geometric matching that contains neither two parallel edges nor a cross-blocker.*

Consider a generalized wheel GW_N with $2n$ vertices and interior vertex v_0. Let us try to construct a spine matching on GW_N. To this end, we first connect v_0 to some other point v_1 with an edge $e = v_0v_1$. Note that the remaining points are now in convex position and hence, there is a unique perfect matching on them without parallel edges, namely, the matching consisting of the halving edges of $GW_N \setminus \{v_0, v_1\}$. Indeed, if a perfect matching on a convex point set contains a non-halving edge, the side of larger cardinality must determine a parallel edge. Thus, for each choice of e we get a unique possible matching, which we call a *potential matching*, and this matching is a spine matching unless some edge is parallel to e or there are two edges that together with e form a cross-blocker. In the following, we investigate the conditions, under which these cases occur.

Consider a non-radial halving edge h of GW_N. Then we call $\mathrm{cl}(h^-)$ a *bad halfplane*. Note that there might be no bad halfplanes, for example if GW_N is a regular wheel. Let us emphasize again that bad halfplanes are closed and in particular, due to the properties of halving edges, the intersection of bad halfplanes either is empty or contains a vertex of GW_N.

Lemma 5.6. *Let GW_N be a generalized wheel and assume it has a bad halfplane B bounded by an edge h. Assume M is a spine matching on GW_N which contains the edge $e = v_0v_1$. Then the vertex v_1 lies in the bad halfplane B.*

Proof. Assume for the sake of contradiction that v_1 does not lie in B. Then, $GW_N \setminus \{v_0, v_1\}$ contains two more vertices in B than the other (closed) side of h. Let f and g be the two edges in M incident to the endpoints of h. Since

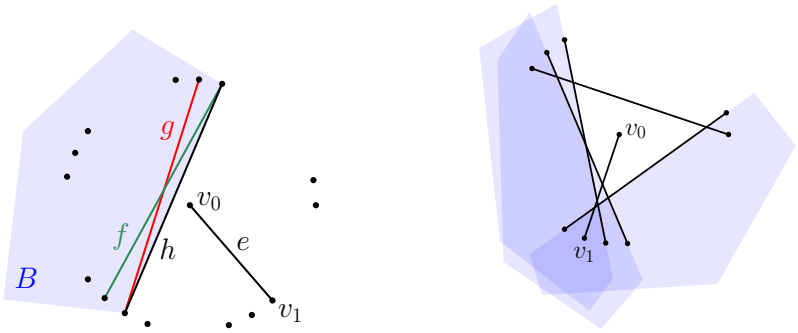

Figure 5.6: (a) Illustration of the proof of Lemma 5.6. The edges e, f, g form a cross-blocker. (b) The drawn edges form a spine matching M and every bad halfplane (blue) contains the vertex v_1 (Lemma 5.6). Furthermore, M is also the potential matching defined by the halving edge v_0v_1 and it does not contain any parallel edges (Lemma 5.8).

M is a spine matching and by the arguments following Corollary 5.5, f and g need to be halving edges in $GW_N \setminus \{v_0, v_1\}$. Hence, f and g connect to vertices in B. But then h separates f and g from e, and thus either one of them is parallel to e or the three of them form a cross-blocker; see Figure 5.6(a). A contradiction to M being a spine matching. □

See also Figure 5.6(b) for an illustration of Lemma 5.6. We immediately derive the following corollary:

Corollary 5.7. *If GW_N has a collection of bad halfplanes whose intersection is empty, then GW_N cannot be partitioned into plane spanning double stars.*

Proof. By Lemma 5.6, v_1 (the vertex connected to v_0) lies in the intersection of all bad halfplanes, implying that the intersection of all bad halfplanes is non-empty. □

Next, we prove that the other direction of Corollary 5.7 holds as well. To this end, we first derive some preliminary lemmas (again, see Figure 5.6(b) for an illustration of Lemma 5.8).

Lemma 5.8. *Let $e = v_0v_1$ be a halving edge of GW_N. Then the potential matching defined by e has no parallel edges.*

Proof. By the construction of the potential matching, for every pair of parallel edges, one of the edges must be e. Assume for the sake of contradiction that there is an edge f parallel to e. Since every edge of the potential matching is a halving edge in $GW_N \setminus \{v_0, v_1\}$, f contains $n-2$ vertices on either side (recall that GW_N contains $2n$ vertices in total). Furthermore, since e and f are

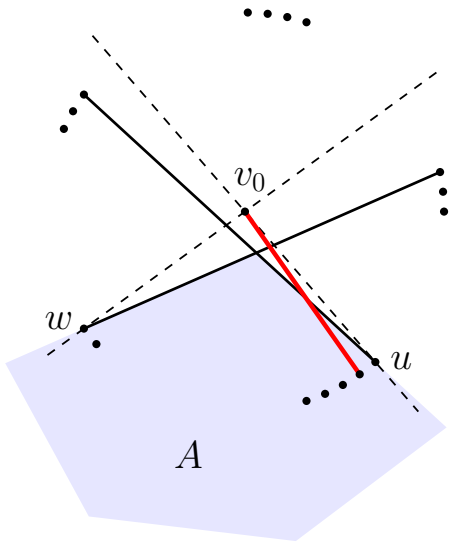

Figure 5.7: The solid edges are halving edges defining bad halfplanes with intersection A. The dashed lines contain at least n vertices on the side containing A. Rotating a line from v_0u to v_0w will therefore find a halving edge with vertex in A (red).

parallel, one side of e contains both endpoints of f and the $n-2$ vertices from one side of f. Hence, e contains n vertices on the same side; a contradiction to e being a halving edge. □

For the following lemma we use a standard rotation argument.

Lemma 5.9. *Let A be a non-empty intersection of bad halfplanes. Then A contains a vertex v_i such that v_0v_i is a halving edge of GW_N.*

Proof. First note that the intersection of the bad halfplanes contains a vertex of GW_N (recall that bad halfplanes are closed). Let u and w be the first and the last vertex of GW_N in this intersection (in clockwise order). Note that u and w are both incident to one of the respective halving edges each. Since bad halfplanes do not contain the center v_0, the corresponding lines through the radial edges v_0u and v_0w have at least n vertices on the side containing A (see Figure 5.7). Rotating a line ℓ through v_0 from u to w (alongside A) will therefore yield a radial halving edge of GW_N with endpoint in A (because hitting a vertex in A decreases the number of vertices on the "right" side of ℓ, while hitting a vertex on the opposite side of A increases that value). □

We are now ready to prove the other direction of Corollary 5.7:

Lemma 5.10. *If GW_N cannot be partitioned into plane spanning double stars, then GW_N has three bad halfplanes whose intersection is empty.*

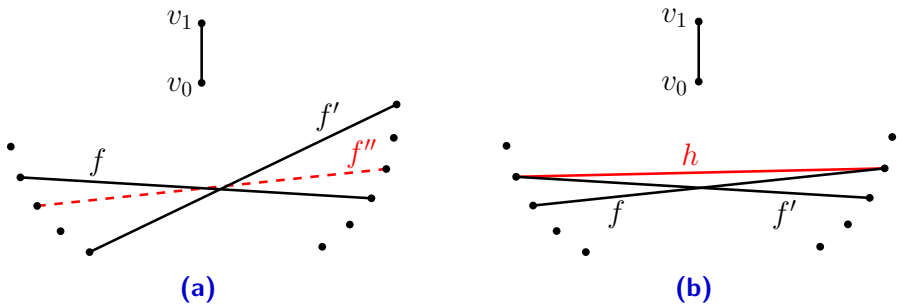

Figure 5.8: Illustration of the proof of Lemma 5.10. (a) Replacing f' by f'' yields a cross blocker with the consecutive diagonal edges f and f''. (b) Since f and f' are halving edges in $GW_N \setminus \{v_0, v_1\}$, the edge h is a halving edge in GW_N.

Proof. Assume that GW_N cannot be partitioned into plane double stars, that is, for every radial edge e the resulting potential matching contains either an edge parallel to e or two edges which together with e form a cross-blocker.

Since every vertex is incident to at least one halving edge (again using a standard rotation argument), we can now consider the potential matching M defined by some radial halving edge $e = v_0v_1$. By Lemma 5.8, M has no parallel edges. Since M is not a spine matching by assumption and it does not contain any parallel edges, it has to contain a cross blocker. Let $f = uw$ and $f' = u'w'$ be the diagonal edges in this cross blocker. We pick a cross blocker in such a way that u and u' as well as w and w' are consecutive along the convex hull (see Figure 5.8(a) for an illustration). Then the unique minimal edge h such that $f <_c h$ and $f' <_c h$ is a halving edge in GW_N; see Figure 5.8(b). Hence, there exists a bad halfplane, which does not contain v_1.

We claim that the intersection of all bad halfplanes is empty. Indeed, if it was not empty, then by Lemma 5.9 the intersection would contain a point v_i such that v_0v_i is a halving edge. But then, by the above arguments, there is a bad halfplane which does not contain v_i, which is a contradiction to v_i lying in the intersection of all bad halfplanes.

As the halfplanes are convex, it now follows from Helly's theorem that if the whole family has empty intersection, then there are some three bad halfplanes whose intersection is already empty. Also note that there cannot be two bad halfplanes with empty intersection. □

To summarize, we get:

Corollary 5.11. *A generalized wheel GW_N cannot be partitioned into plane spanning double stars if and only if GW_N has three bad halfplanes whose intersection is empty.*

Proof. Follows from Corollary 5.7 and Lemma 5.10. □

Finally, we are ready to prove Theorem 2.4, which we restate here for easier readability:

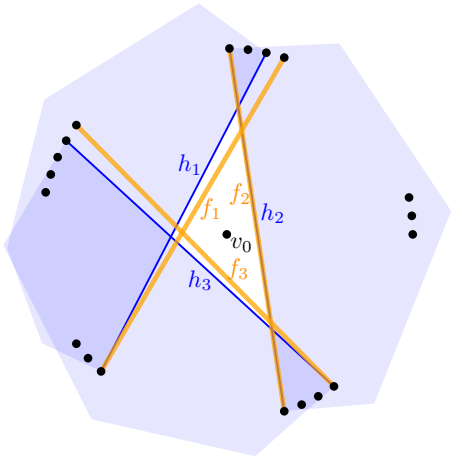

Figure 5.9: Illustration of the proof of Theorem 2.4.

Theorem 2.4. *Let GW_N be a generalized wheel with k groups and $2n$ vertices. Then GW_N cannot be partitioned into plane spanning double stars if and only if there are three families of $\frac{k-1}{2}$ consecutive groups, such that (i) every family contains at most $n-2$ vertices and (ii) every group is in at least one family.*

Proof. Let GW_N be a generalized wheel with k groups and $2n$ vertices. For the proof it is more convenient to consider everything from the side of the (complementary) $\frac{k+1}{2}$ consecutive groups. That is, by Corollary 5.11, it is enough to show that GW_N contains three bad halfplanes whose intersection is empty, if and only if there are three families of $\frac{k+1}{2}$ consecutive groups, each containing at least $n+1$ vertices, such that no group is in all three families.

For the one direction, assume there are three bad halfplanes whose intersection is empty (see Figure 5.9 for an illustration). Let h_1, h_2, h_3 be the three respective halving edges. Next, consider for each of the three halving edges, a maximal diagonal edge f_i (of distance d_1) with $h_i <_c f_i$ (for $i = 1, 2, 3$). Clearly, the closure of each f_i^- contains $\frac{k+1}{2}$ consecutive groups and at least $n+1$ vertices. It remains to show that there is no group contained in all f_i^-. To this end, note first that any pair of bad halfplanes overlaps, that is, contains vertices of GW_N in their intersection. Therefore, the union of the three bad halfplanes covers the entire convex hull of GW_N, which also holds for the union of the three $\mathrm{cl}(f_i^-)$. Assume for the contrary that there is a common group in the intersection of the three maximal diagonal edges, that is, $\bigcap \mathrm{cl}(f_i^-)$ is non-empty. Let x be a point in $\bigcap \mathrm{cl}(f_i^-)$ and x' be the antipodal point. Then x' lies in f_i^+ for any i and hence the convex hull is not fully covered, a contradiction.

For the other direction, let F_1, F_2, F_3 be three families of $\frac{k+1}{2}$ consecutive groups each containing at least $n+1$ vertices such that no group is in all three families. Also let f_1, f_2, f_3 be the (maximal) diagonal edges bounding these

families. Then each $\mathrm{cl}(f_i^-)$ contains a halving edge e_i. It remains to show that the intersection of the corresponding bad halfplanes is empty. Note that a non-empty intersection of three bad halfplanes must also contain a vertex of GW_N. Then, the corresponding group would be a common group of the three F_i's. □

Chapter 6

Beyond-planar Partitions

Given the negative answers to Question 3.1 and Question 4.1, it is natural to study partitions into beyond-planar subgraphs, that is, subgraphs in which some crossings are allowed. Beyond-planar graphs are a natural generalization of planar graphs and a rapidly growing research area. We refer the interested reader to the book of Hong and Tokuyama [Hon20] or the survey of Didimo et al. [DLM19] and references therein.

However, to the best of our knowledge, edge partitions into beyond-planar graphs have not been studied. We initiate this study for two important classes of beyond planar-graphs, namely, *k-plane subgraphs* (where every edge is crossed by at most k other edges) and *k-quasi-plane subgraphs* (in which no k edges pairwise cross). For the former, we show bounds on the number of subgraphs required for partitioning $K(S)$ for S in convex position (Theorem 2.5 and Theorem 2.6). For the latter, we show that a partition into 3-quasi-plane spanning trees is possible for any S with $|S|$ even (Theorem 2.8). This is best possible, as 2-quasi-plane graphs are plane. We further present bounds on the partition of any $K(S)$ into k-quasi-plane subgraphs for general k (Theorem 2.9).

Contribution. In the setting of k-plane partitions we focus on the convex setting, showing the following bounds (which are tight for $k = 1$):

Theorem 2.5. *For a point set S in convex position with $|S| = n \geq 5$, $K(S)$ can be partitioned into $\lceil \frac{n}{3} \rceil$ many 1-plane subgraphs and $\lceil \frac{n}{3} \rceil$ subgraphs are required in every 1-plane partition.*

Theorem 2.6. *For an n-point set S in convex position and every $k \in \mathbb{N}$, $K(S)$ admits a partition into at most $\frac{n}{\sqrt{2k}}$ k-plane subgraphs. More precisely, for every $s > 2$, $K(S)$ admits a $\frac{(s-1)(s-2)}{2}$-plane partition into $\lceil \frac{n}{s} \rceil$ subgraphs. Conversely, for every $k \in \mathbb{N}$, at least $\frac{n-1}{4.93\sqrt{k}}$ subgraphs are required in any k-plane partition of $K(S)$.*

Towards the lower bound of Theorem 2.6, it will be crucial to understand how many edges a k-planar graph can maximally have: Pach and Tóth [PT97]

showed an upper of $\left(4.108\sqrt{k}\cdot n\right)$ in general and an improved bound of $(k+3)(n-2)$ for $k \leq 4$. In our convex setting, we improve these bounds to $\left(2.465\sqrt{k}\cdot n\right)$ and to $\left(\frac{k+4}{2}n-(k+3)\right) = (\frac{k}{2}+2)(n-2)+1$ for $k \leq 4$ (Proposition 6.3 and Proposition 6.1).

Along the way, we also study the well-known crossing lemma, which gives a lower bound on the crossing number of a graph G in terms of $|V(G)|$ and $|E(G)|$. The *crossing number* $\mathrm{cr}(G)$ of a graph G is the minimum number of crossings in any drawing of G. Ajtai et al. [ACNS82] and, independently, Leighton [Lei83] showed that there exists a constant $c > 0$ such that for every graph G with n vertices and $e \geq 4n$ edges, $\mathrm{cr}(G) \geq c \cdot \frac{e^3}{n^2}$ holds. Asymptotically this bound is known to be tight and currently the best constant $c = \frac{1}{29}$ is due to Ackerman [Ack19]. In our convex setting, we derive the following improved bound:

Lemma 2.7 (convex crossing lemma)**.** *Let G be a graph with n vertices and e edges such that $e \geq \frac{9}{2}n$. Then every straight-line drawing of G in which the vertices of G are placed in convex position has at least*

$$\frac{20}{243}\frac{e^3}{n^2} \approx 0.0823\frac{e^3}{n^2}$$

crossings.

Moreover, we consider partitions into k-quasi-plane subgraphs for arbitrary point sets (in general position). We show that a partition into 3-quasi-plane spanning trees is possible for any S with $|S|$ even. This is best possible, as 2-quasi-plane graphs are plane. We further present bounds on the partition of any $K(S)$ into k-quasi-plane subgraphs for general k.

Theorem 2.8. *Let S be a point set of size $2n$, then the complete geometric graph $K(S)$ can be partitioned into n 3-quasi-plane spanning trees.*

Theorem 2.9. *Let S be a set of n points in general position and denote the size of a largest crossing family on S by m. Also let $k \in \mathbb{N}$ such that $3 \leq k \leq m$. Then, at least $\lceil \frac{m}{k-1} \rceil$ subgraphs are required and at most $\lceil \frac{m}{k-1} \rceil + \lceil \frac{n-2m}{k-1} \rceil$ subgraphs are needed to partition the complete geometric graph $K(S)$ into k-quasi-plane subgraphs.*

Organization of this chapter. Section 6.1 and Section 6.2 are dedicated to the k-plane setting, while Section 6.3 is concerned with k-quasi-plane subgraphs. In Section 6.1 we focus on the case $k = 1$ and prove Theorem 2.5. In Section 6.2 we generalize to arbitrary k, proving Theorem 2.6 and Lemma 2.7. Finally, in Section 6.3 we prove Theorem 2.8 and Theorem 2.9.

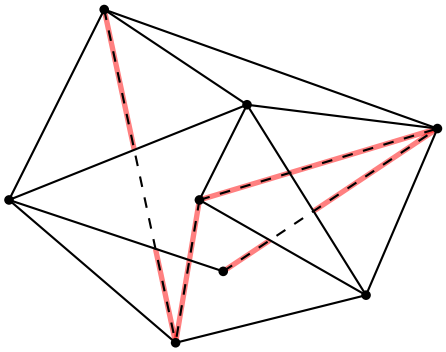

Figure 6.1: A 2-plane graph M (black solid and dashed) and a plane subgraph M' (solid black) with a maximum number of edges. The eight partial edges are depicted in red.

6.1 Partitions into 1-plane subgraphs

In this section, we prove Theorem 2.5. As the upper bound follows immediately from Theorem 2.6, which we prove in the following section, we are only concerned with the lower bound. To this end, we first prove a proposition concerning the number of edges that a k-planar graph can maximally have for small values of k:

Proposition 6.1. *For every $k \leq 4$, every straight-line convex k-plane graph G on $n \geq 2$ vertices has at most $\left(\frac{k+4}{2}n - (k+3)\right)$ edges.*

For the proof we need the following notion: Given a k-plane graph M, consider a plane subgraph M' of M with a maximum number of edges. Then every edge $e \in E(M) \setminus E(M')$ must cross at least one edge of M'. For $e \in E(M) \setminus E(M')$, the closed portion between an endpoint of e and the nearest crossing of e with an edge of M' is called a *partial edge*. Note that every edge in $E(M) \setminus E(M')$ contributes exactly two partial edges. Furthermore, every partial edge is fully contained in a face of M'. See Figure 6.1 for an illustration.

Pach and Tóth [PT97] proved the following bound concerning the maximal number of partial edges[1]:

Lemma 6.2 (cf. [PT97, Lemma 2.1])**.** *Let $k \leq 4$ and let M be a k-plane graph. Let M' be a plane subgraph of M with a maximum number of edges. Let Φ be a face in M' with $|\Phi| \geq 3$, whose boundary is connected. Then there are at most $(|\Phi| - 2)(k+1) - 1$ partial edges contained in the closed interior of Φ.*

With this lemma at hand, we are now ready to prove Proposition 6.1. Recall that $|\Phi|$ counts the number of bounding edges of Φ with multiplicities.

[1]Note that Pach and Tóth use the term *half-edge*. However, as this term is also commonly used in other contexts, we decided to use partial edge.

Proof of Proposition 6.1. Let $k \leq 4$ and let G be a straight-line, convex k-plane graph. Furthermore, let G' be a plane subgraph of G with a maximum number of edges. W.l.o.g. we may assume that G and hence, also G' contains all n convex hull edges.

Thus every face Φ in G' has a connected boundary and hence, using Lemma 6.2, contains at most $(|\Phi| - 2)(k+1) - 1$ partial edges. Further note that the unbounded face does not contain any partial edge. Denote the set of bounded faces of G' by $\mathcal{F}$. Then, using the fact that the unbounded face of G' contributes exactly n edges, it follows:

$$\begin{aligned}\sum_{\Phi \in \mathcal{F}} (|\Phi| - 2) &= (2e(G') - n) - 2|\mathcal{F}| \\ &= (2e(G') - n) - 2(f(G') - 1) \\ &= 2(e(G') - n - f(G')) + n + 2 \\ &= n - 2,\end{aligned}$$

where in the last equality we used Euler's formula $n - e(G') + f(G') = 2$.

Recall that every partial edge is contained in a bounded face of G' and hence, using Lemma 6.2, the total number of partial edges is at most

$$\begin{aligned}\sum_{\Phi \in \mathcal{F}} ((|\Phi| - 2)(k+1) - 1) &= k \sum_{\Phi \in \mathcal{F}} (|\Phi| - 2) + \sum_{\Phi \in \mathcal{F}} |\Phi| - 3|\mathcal{F}| \\ &= k(n-2) + (2e(G') - n) - 3(f(G') - 1).\end{aligned}$$

On the other hand, since every edge in $E(G) \setminus E(G')$ contributes precisely two partial edges, the total number of partial edges equals $2(e(G) - e(G'))$, and hence we conclude:

$$2e(G) - 2e(G') \leq k(n-2) + (2e(G') - n) - 3f(G') + 3.$$

Rearranging, again applying Euler's formula $n - e(G') + f(G') = 2$, and using the fact that $f(G') \leq n - 1$ holds, now yields:

$$\begin{aligned}2e(G) &\leq k(n-2) - n + 3 + 4e(G') - 3f(G') \\ &= k(n-2) - n + 3 + 4(e(G') - f(G')) + f(G') \\ &\leq k(n-2) - n + 3 + 4(n-2) + (n-1) \\ &= (k+4)n - (2k+6).\end{aligned}$$

Dividing by two yields the claimed bound $e(G) \leq \frac{k+4}{2}n - (k+3)$, as desired. □

Next, we prove Theorem 2.5.

Theorem 2.5. *For a point set S in convex position with $|S| = n \geq 5$, $K(S)$ can be partitioned into $\lceil \frac{n}{3} \rceil$ many 1-plane subgraphs and $\lceil \frac{n}{3} \rceil$ subgraphs are required in every 1-plane partition.*

Proof. The upper bound follows from Theorem 2.6 with $s = 3$.

For the lower bound, suppose we are given a 1-plane partition $E_1, \dots, E_c$ of $K(S)$. Then we show that $c \geq \lceil \frac{n}{3} \rceil$ holds. Let $O \subseteq E(K(S))$ be the set of the n boundary edges of $K(S)$, and note that they are not crossed by any other edge. Hence, for every $i \in \{1, \dots, c\}$, the geometric graph formed by the points in S and the edges in $O \cup E_i$ is 1-plane, and hence by Proposition 6.1 we have $|O \cup E_i| \leq \frac{5}{2}n - 4$ for all $i \in \{1, \dots, c\}$. This clearly implies that $|E_i \setminus O| \leq \left(\frac{5}{2}n - 4\right) - n = \frac{3}{2}n - 4$ for every i. Since the sets $\{E_i \setminus O\}_{i \in \{1,\dots,c\}}$ form a partition of the set of all $\frac{n(n-1)}{2} - n = \frac{n(n-3)}{2}$ interior edges of $K(S)$, we conclude:

$$c \geq \frac{\left(\frac{n(n-3)}{2}\right)}{\frac{3}{2}n - 4} = \frac{n(n-3)}{3n-8} = \frac{n}{3} - \frac{n}{3(3n-8)} > \frac{n}{3} - \frac{1}{3}.$$

(For the last inequality we used $n \geq 5$). This implies $c \geq \lceil \frac{n}{3} \rceil$, as claimed. This concludes the proof. □

6.2 Partitions into k-plane subgraphs

For the proof of Theorem 2.6, we need to generalize Proposition 6.1 to larger values of k. To this end, we first prove the following convex crossing lemma. We adapt the standard probabilistic proof from the textbooks (see e.g. [AZ99]).

Lemma 2.7 (convex crossing lemma)**.** *Let G be a graph with n vertices and e edges such that $e \geq \frac{9}{2}n$. Then every straight-line drawing of G in which the vertices of G are placed in convex position has at least*

$$\frac{20}{243}\frac{e^3}{n^2} \approx 0.0823\frac{e^3}{n^2}$$

crossings.

Proof. We consider the following process: We start with G, and repeatedly update a subgraph of G. As long as the current subgraph still contains crossings, pick an edge with the highest number of crossings from the current drawing and remove it. Repeat until we end up with a crossing-free subgraph. As long as the current subgraph has more than $4n - 7$ edges, by Proposition 6.1 applied for $k = 4$, it follows that the current subgraph contains edges which cross with at least 5 other edges, and hence, in the next step the edge we remove will remove at least 5 crossings from G. Similarly, as long as the number of edges in the subgraph is strictly greater than $\frac{7}{2}n - 6$, by Proposition 6.1 the edges we remove will be edges with at least 4 crossings, if it is strictly greater than $3n - 5$, then we remove at least 3 crossings at each step, etc.

Summing up all these different contributions of how many crossings are lost in the process by removing edges, we obtain the following lower bound on the number of crossings of G:

$$\operatorname{cr}(G) \geq 5\left(e-(4n-7)\right)+4\left((4n-7)-\left(\frac{7}{2}n-6\right)\right)+3\left(\left(\frac{7}{2}n-6\right)-(3n-5)\right)+$$
$$2\left((3n-5)-\left(\frac{5}{2}n-4\right)\right)+\left(\left(\frac{5}{2}n-4\right)-(2n-3)\right)=5e-15n+25.$$

So, let $p \in [0,1]$ be a probability to be determined later and let G_p be the induced random subgraph of G, where each vertex is picked with probability p. Then we have:

$$\mathbb{E}[e(G_p)] = p^2 e(G)$$
$$\mathbb{E}[v(G_p)] = pn$$
$$\mathbb{E}[\operatorname{cr}(G_p)] = p^4 \cdot \operatorname{cr}(G)$$

By the above inequality, which also holds for the convex geometric subgraph G_p of G, we get:

$$\operatorname{cr}(G_p) \geq 5e(G_p) - 15v(G_p)$$

Taking expectations and dividing by p^4, yields:

$$\operatorname{cr}(G) \geq \frac{5e(G)}{p^2} - 15\frac{n}{p^3}$$

Plugging in the optimal value $p^* = \frac{9n}{2e} \leq 1$ (here we used our assumption $e \geq \frac{9}{2}n$) yields the desired result. □

Using this lemma, we can generalize Proposition 6.1 as follows:

Proposition 6.3. *For every $k \geq 5$, every convex k-plane graph G on n vertices has at most $\sqrt{\frac{243}{40}k} \cdot n \approx 2.465\sqrt{k} \cdot n$ edges.*

Proof. Let $k \geq 5$ and let G be a convex k-plane graph with n vertices. If $e(G) \leq \frac{9}{2}n$, there is nothing to show, because $\sqrt{\frac{243}{40}k} > \frac{9}{2}$ holds for every $k \geq 5$. Hence, we assume $e := e(G) > \frac{9}{2}n$ and apply the convex crossing lemma to conclude:

$$\operatorname{cr}(G) \geq \frac{20}{243}\frac{e^3}{n^2}.$$

On the other hand, since G is convex k-plane we know that every edge participates in at most k crossings. By double-counting (noting that exactly two edges participate in every crossing), we get that

$$\frac{ke}{2} \geq \operatorname{cr}(G) \geq \frac{20}{243}\frac{e^3}{n^2}.$$

Rearranging the inequality yields the claim. □

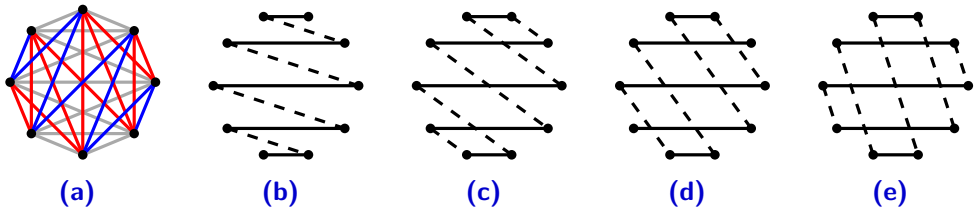

Figure 6.2: (a) Partition into 1-plane subgraphs by composing groups of (at most) 3 consecutive slopes each. (b)-(e) Edges with slope distance 1/2/3/4 intersect at most 0/1/2/3 times.

Finally, we are ready to prove Theorem 2.6. Also recall that we associate subgraphs with colors for convenience.

Theorem 2.6. *For an n-point set S in convex position and every $k \in \mathbb{N}$, $K(S)$ admits a partition into at most $\frac{n}{\sqrt{2k}}$ k-plane subgraphs. More precisely, for every $s > 2$, $K(S)$ admits a $\frac{(s-1)(s-2)}{2}$-plane partition into $\lceil \frac{n}{s} \rceil$ subgraphs. Conversely, for every $k \in \mathbb{N}$, at least $\frac{n-1}{4.93\sqrt{k}}$ subgraphs are required in any k-plane partition of $K(S)$.*

Proof. Let us first prove the upper bound. To this end, suppose that $s \geq 2$ is such that $\frac{(s-1)(s-2)}{2} \leq k$, and let us show that $K(S)$ can be partitioned into $\lceil \frac{n}{s} \rceil$ k-plane subgraphs. W.l.o.g. assume that the points in S form a regular n-gon (as the crossings are determined by the rotation system). Then every edge defines a slope and in total there are n different slope values, which we sort in circular order. Next, we partition this list of slope values into $\lceil \frac{n}{s} \rceil$ (contiguous) intervals of size at most s. Then, define a color class for all edges whose slopes fall into a common interval of this partition, see Figure 6.2(a).

We show that all these subgraphs are $\frac{(s-1)(s-2)}{2}$-plane. To this end, define the *slope distance* to be the distance between two slope values in the circularly sorted list of slopes. Note that edges cannot be crossed by other edges of the same slope or slope distance 1; by at most one edge of slope distance 2, by at most two edges of slope distance 3, etc. (see Figures 6.2(b) to 6.2(e)). Hence, if an edge e has color i, and if the slope of e is the j-th slope ($j \in \{1, \ldots, s\}$) in its circular interval of slopes, then e can cross with at most the following amount of edges of color i:

$$\sum_{1 \leq k < j-1} (j-k-1) + \sum_{j+1 < k \leq s} (k-j-1) = \frac{(j-1)(j-2)}{2} + \frac{(s-j)(s-j-1)}{2} =$$
$$= \frac{(s-1)(s-2)}{2} - (s-j)(j-1) \leq \frac{(s-1)(s-2)}{2}.$$

For the lower bound, note that $K(S)$ has $\frac{n(n-1)}{2}$ edges, and that in every k-plane partition of $K(S)$, every color class induces a convex k-plane subgraph

on n vertices. Hence, by Proposition 6.3, every color class has size at most $\sqrt{\frac{243}{40}k} \cdot n$. So, the number of colors required in any k-plane partition is at least

$$\frac{\binom{n(n-1)}{2}}{\sqrt{\frac{243}{40}k} \cdot n} \geq \frac{n-1}{4.93\sqrt{k}}.$$

This concludes the proof. □

The following intriguing question is left open by our study.

Question 6.4. *Is the upper bound in Theorem 2.6 tight up to lower-order terms?*

More generally, it would be interesting to shed some more light on the "in-between-cases" coming out of the upper bound in Theorem 2.6: the term $\frac{(s-1)(s-2)}{2}$ covers only the values $0, 1, 3, 6, 10, \ldots$. For instance, can we partition convex complete geometric graphs with fewer colors into 2-plane subgraphs than we need for the 1-plane partition? More generally, for $\frac{(s-1)(s-2)}{2} < k < \frac{s(s-1)}{2}$, can we improve upon the $\lceil \frac{n}{s} \rceil$ bound from Theorem 2.6 for k-plane partitions? This question is surprisingly difficult (even for $k = 2$)[2] and we do not know of any improvements of the bounds for these "in-between-cases".

6.3 Partitions into k-quasi-plane subgraphs and spanning trees

In this section, we develop bounds on the number of colors required in a k-quasi-plane partition for point sets in general position (for $k = 2$ this again amounts to the setting of plane subgraphs and hence, we assume $k \geq 3$ in the following). The setting of spanning trees is resolved by Theorem 2.8:

Theorem 2.8. *Let S be a point set of size $2n$, then the complete geometric graph $K(S)$ can be partitioned into n 3-quasi-plane spanning trees.*

Proof. Let $p_1, \ldots, p_{2n}$ be the points of S sorted by x-coordinate (w.l.o.g. no two points have the same x-coordinate). We distinguish the points into evenly indexed points p_{2i} for $i = 1, \ldots, n$ and oddly indexed points p_{2i-1} for $i = 1, \ldots, n$.

The goal is to define n double stars that partition $K(S)$. We construct a double star T_i for every two consecutive vertices $\{p_{2i-1}, p_{2i}\}$, for $i = 1, \ldots, n$ as follows (see Figure 6.3 for an illustration). We connect p_{2i-1} with all evenly indexed points left of p_{2i-1} and all oddly indexed points right of p_{2i-1}. Further, we connect p_{2i} with the remaining vertices, namely with all oddly indexed

[2] Using computer assistance, we can show that at least $\frac{3n}{10}$ colors are required for any 2-plane partition, almost matching the $\frac{n}{3}$ bound from the 1-plane partition.

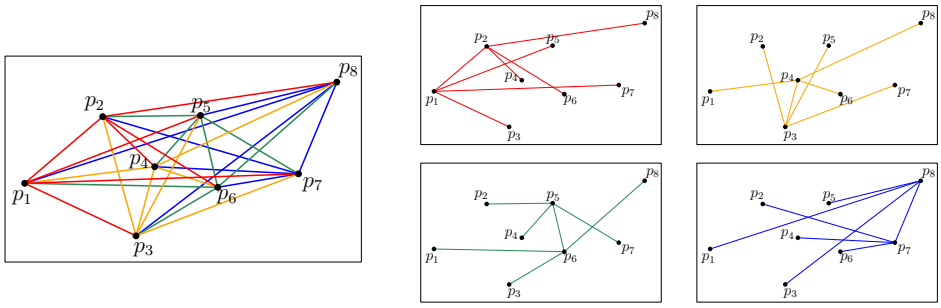

Figure 6.3: Illustration of Theorem 2.8.

points left of p_{2i} and all evenly indexed points right of p_{2i}. Note that this includes the edge $p_{2i-1}p_{2i}$. T_i indeed forms a double star, as all vertices are connected to either p_{2i-1} or p_{2i}. In particular T_i forms a spanning tree.

It remains to show that every edge is covered exactly once. Let $e = p_ip_j$ be an edge $(1 \leq i < j \leq n)$. If i and j have the same parity, e belongs to T_i and otherwise to T_j.

Since every double star is necessarily 3-quasi-plane, the claim follows. □

So, we turn our attention to the subgraph setting. Consider a point set S of size $2n$ with a crossing family of size n. Then all edges involved in this crossing family are halving edges. Let $\ell_1, \ldots, \ell_n$ be the corresponding halving lines. We label these lines in clockwise order, more precisely their intersections with a sufficiently large circle appear in clockwise order (where it follows from the properties of halving lines that this enumeration is consistent on both sides). For each halving line ℓ_i define an infinitesimally counter-clockwise rotated line ℓ_i', such that the two defining vertices (say p_i, q_i) of ℓ_i lie to either side of ℓ_i'. Define ℓ_i^+ to be the upper halfplane (bounded by ℓ_i') and let it contain p_i; similarly ℓ_i^- denotes the lower halfplane (bounded by ℓ_i') and it contains q_i. See Figure 6.4 for an illustration.

Theorem 6.5. *Let S be a point set of size $2n$ with a crossing family of size n and let $k \geq 3$. Then $\lceil \frac{n}{k-1} \rceil$ colors are required and sufficient to partition $K(S)$ into k-quasi-plane subgraphs.*

Proof. The lower bound is easy: since we have a crossing family of size n, we need at least $\lceil \frac{n}{k-1} \rceil$ different colors to partition $K(S)$ into k-quasi-plane subgraphs.

The other direction is more involved. We build $c := \lceil \frac{n}{k-1} \rceil$ subgraphs $G_1, \ldots, G_c$ of $K(S)$ as follows. Each subgraph G_i in turn is formed by the union of three subgraphs.

To construct G_1, let X_1 be the collection of vertices of the first $k-1$ consecutive halving lines starting from ℓ_1, i.e., $X_1 := \{p_1, \ldots, p_{k-1}, q_1, \ldots, q_{k-1}\}$. Next, we consider all points in ℓ_1^+ and form the complete bipartite graph B_1^+

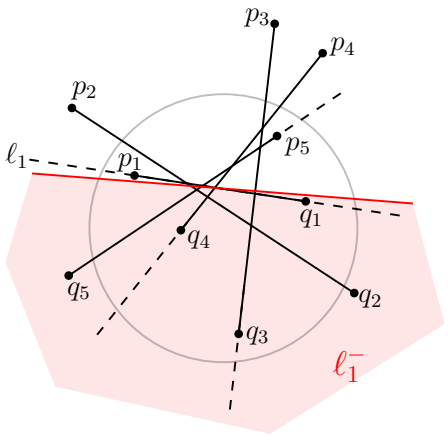

Figure 6.4: The labeling of the halving lines $\ell_1, \ldots, \ell_n$ in clockwise order. Illustration of the definition of the halfplane ℓ_1^-. Note that $p_i \in \ell_1^+$ and $q_i \in \ell_1^-$ for any i.

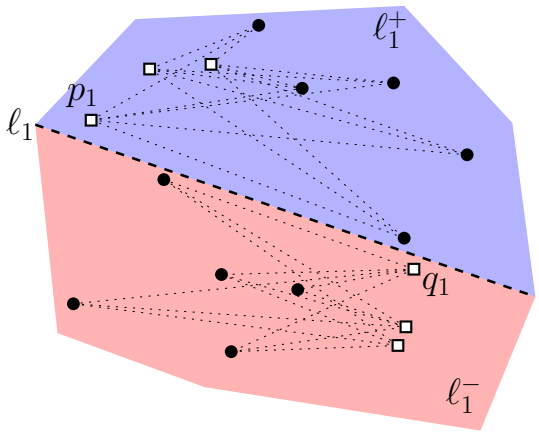

Figure 6.5: The vertices of X_1 are represented by squares. The dotted edges in the blue (red) region represent the complete bipartite subgraph B_1^+ (B_1^-) of G_1 corresponding to the line ℓ_1 with the defining vertices p_1 and q_1.

between the points in $(X_1 \cap \ell_1^+)$ and those in $((S \setminus X_1) \cap \ell_1^+)$. Symmetrically, we form the complete bipartite graph B_1^- between the point set $X_1 \cap \ell_1^-$ and the point set $(P \setminus X_1) \cap \ell_1^-$. An illustration of the construction is given in Figure 6.5. The subgraph G_1 is finally defined to be the union of $K(X_1), B_1^+$ and B_1^-.

We iteratively repeat the same process for the next $k-1$ halving lines until reaching p_n. More precisely, G_l consists of the union of the complete graph with vertex set $X_l = \{p_i, q_i \,|\, (l-1)(k-1) < i \leq \min(l(k-1), n)\}$ and the two bipartite graphs defined by $\ell_{(l-1)(k-1)+1}$ as before. The last graph G_c may be formed by less than $k-1$ halving lines.

We first validate that each G_i is k-quasi-plane. For notational convenience we prove it only for G_1 in the following, though the same argument works for any G_i. The two bipartite subgraphs B_1^+ and B_1^- are disjoint, as they lie to either side of the line ℓ_1'. Thus, any crossing family contains only edges from at most one of them, and potentially further edges from the complete subgraph $K(X_1)$. Therefore every edge from the crossing family must be incident to one of the vertices that are both in X_1 and on the respective side of ℓ_1. By construction, there can be at most $k-1$ such vertices and therefore, any crossing family has at most size $k-1$ as well.

It remains to show that any edge of $K(S)$ is covered by some G_i. Let $e := \{u, v\}$ be an edge of $K(S)$. For every $i = 1, \ldots, n$, p_i and q_i are part of $X_{\lceil i/(k-1)\rceil}$. Thus, the endpoint u is contained in X_r for some r, and v is contained in X_s, for some s. If $r = s$ we are done, as then e is contained in $K(X_r)$ and thus part of G_r. So suppose $r \neq s$. Then v and u lie on the same side of either ℓ_r' or ℓ_s'. In the former case, e would be contained in either of the two complete bipartite subgraphs of G_r, that is, B_r^+ or B_r^-. In the latter, e is contained in G_s. □

The restriction in Theorem 6.5 to have a crossing family of size n is rather restrictive, however, not nearly as restricitive as e.g. convex position. Given that every point set determines a near linear size crossing family [PRT19], it is natural to ask whether Theorem 6.5 can be adjusted to point sets in general position. We leave this as an interesting open question (cf. Question 10.5).

Theorem 2.9. *Let S be a set of n points in general position and denote the size of a largest crossing family on S by m. Also let $k \in \mathbb{N}$ such that $3 \leq k \leq m$. Then, at least $\lceil \frac{m}{k-1} \rceil$ subgraphs are required and at most $\lceil \frac{m}{k-1} \rceil + \lceil \frac{n-2m}{k-1} \rceil$ subgraphs are needed to partition the complete geometric graph $K(S)$ into k-quasi-plane subgraphs.*

Proof. Let $S' \subseteq S$ be the subset of endpoints induced by a largest crossing family of size m.

Then, the lower bound follows immediately from Theorem 6.5 applied on S'.

For the upper bound, divide the point set $S \setminus S'$ into disjoint subsets $Q_1, \ldots, Q_c$ of size $k-1$, where $c = \lceil \frac{n-2m}{k-1} \rceil$. For each edge with an endpoint in some Q_i assign it the color i (for edges that have two choices, pick one arbitrarily). Certainly, each color class is k-quasi-plane, since it consists of (at most) the union of $k-1$ stars. Together with $K(S')$, which we can partition by using $\lceil \frac{m}{k-1} \rceil$ colors, the upper bound follows. □

Part II

Flip Graphs

Chapter 7

Flipping Plane Spanning Paths

The focus of this chapter is the following conjecture posed by Akl et al. [AIM07]:

Conjecture 7.1 (Akl et al. [AIM07]). *For every point set S in general position, the flip graph on $\mathcal{P}(S)$ is connected.*

Recall that $\mathcal{P}(S)$ denotes the set of all plane straight-line spanning paths on S. Also recall that a *flip* on a path $P \in \mathcal{P}(S)$ is the operation of removing an edge e from P and replacing it by another edge f on S such that the graph $(P \setminus e) \cup f$ is again a valid path from $\mathcal{P}(S)$. Note that the edges e and f might cross.

Related work. In the setting of plane spanning paths, flips are rather restricted, making it difficult to prove a positive answer. Prior to our work only results for point sets in convex position and very small point sets were known. Akl et al. [AIM07], who initiated the study of flip connectivity on plane spanning paths, showed connectedness of the flip graph on $\mathcal{P}(S)$ if S is in convex position or $|S| \leq 8$. In the convex setting, Chang and Wu [CW09] derived tight bounds concerning the diameter of the flip graph, namely, $2n-5$ for $n = 3, 4$, and $2n-6$ for $n \geq 5$.

Furthermore, using the order type database [AAK02], Aichholzer [Aic] computationally verified Conjecture 7.1 for every set of $n \leq 10$ points in general position (even when using only Type 1 flips).[1]

For the remainder of this chapter, we consider the flip graph on $\mathcal{P}(S)$ (or a subset of $\mathcal{P}(S)$). Moreover, unless stated otherwise, the word *path* always refers to a path from $\mathcal{P}(S)$ for an underlying point set S that is clear from the context.

Flips in plane spanning paths. Let us have a closer look at the different types of possible flips for a path $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ (see also Figure 7.1).

[1] The source code is available at https://github.com/jogo23/flipping_plane_spanning_paths.

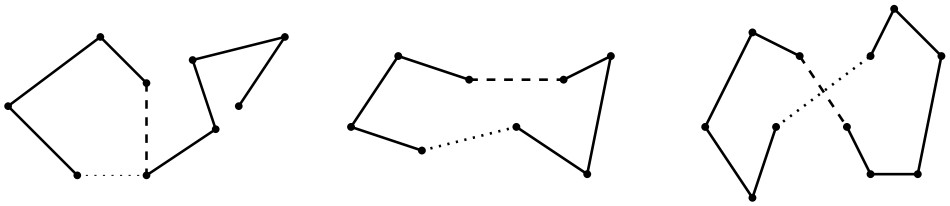

(a) Type 1 flip. **(b)** Type 2 flip. **(c)** Type 3 flip.

Figure 7.1: The three types of flips in plane spanning paths.

When removing an edge $v_{i-1}v_i$ from P with $2 \leq i \leq n$, there are three possible new edges that can be added in order to obtain a path (where, of course, not all three choices will necessarily lead to a valid path in $\mathcal{P}(S)$): v_1v_i, $v_{i-1}v_n$, and v_1v_n. A flip of *Type 1* adds the edge v_1v_i (if $i > 2$) or the edge $v_{i-1}v_n$ (if $i < n$). It results in the path $v_{i-1}, \ldots, v_1, v_i, \ldots, v_n$, or the path $v_1, \ldots, v_{i-1}, v_n, \ldots, v_i$. That is, a Type 1 flip inverts a contiguous chunk from one of the two ends of the path. A flip of *Type 2* adds the edge v_1v_n and has the additional property that the edges $v_{i-1}v_i$ and v_1v_n do not cross. In this case, the path P together with the edge v_1v_n forms a plane cycle. If a Type 2 flip is possible for one edge $v_{i-1}v_i$ of P, then it is possible for all edges of P. A flip of *Type 3* also adds the edge v_1v_n, but now the edges v_1v_n and $v_{i-1}v_i$ cross. Note that a Type 3 flip is only possible if the edge v_1v_n crosses exactly one edge of P, and then the flip is possible only for the edge $v_{i-1}v_i$ that is crossed. We further remark that Type 2 flips are not relevant for the *connectedness* of the flip graph but only for the diameter, since every Type 2 flip can be simulated by a sequence of Type 1 flips, e.g., flip v_1v_2 to v_1v_n, then flip v_2v_3 to v_1v_2, then v_3v_4 to v_2v_3, etc., until flipping $v_{i-1}v_i$ to $v_{i-2}v_{i-1}$. For Type 3 flips it remains open whether they are relevant for the connectedness of the flip graph (in the following, we are not using any Type 3 flip). Unless mentioned otherwise, all flips we are using in the following are Type 1 flips.

Contribution. We approach Conjecture 7.1 from two directions. First, we show that it is sufficient to prove flip connectivity for paths with a fixed starting edge. Second, we verify Conjecture 7.1 for several classes of point sets, namely wheel sets and generalized double circles (which include, e.g., double chains and double circles).

Towards the first part, we define, for two distinct points $p, q \in S$, the following subsets of $\mathcal{P}(S)$: let $\mathcal{P}(S, p)$ be the set of all plane spanning paths for S that start at p, and let $\mathcal{P}(S, p, q)$ be the set of all plane spanning paths for S that start at p and continue with q. Then for any S, the flip graph on $\mathcal{P}(S, p, q)$ is a subgraph of the flip graph on $\mathcal{P}(S, p)$, which in turn is a subgraph of the flip graph on $\mathcal{P}(S)$. We conjecture that all these flip graphs are connected:

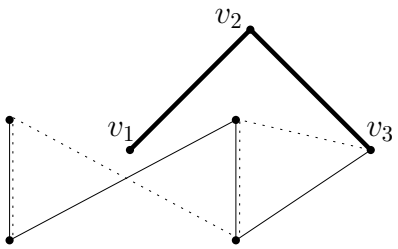

Figure 7.2: Example where the flip graph is disconnected if the first three vertices of the paths are fixed. No edge of the solid path can be flipped, but there is at least one other path (dotted) with the same three starting vertices.

Conjecture 7.2. *For every point set S in general position and every $p \in S$, the flip graph on $\mathcal{P}(S, p)$ is connected.*

Conjecture 7.3. *For every point set S in general position and every $p, q \in S$, the flip graph on $\mathcal{P}(S, p, q)$ is connected.*

Towards Conjecture 7.1, we show that it suffices to prove Conjecture 7.3:

Theorem 7.4. *Conjecture 7.2 implies Conjecture 7.1.*

Theorem 7.5. *Conjecture 7.3 implies Conjecture 7.2.*

Note that the analogue of Conjecture 7.3 for paths where the first $k \geq 3$ vertices are fixed, does not hold: Figure 7.2 shows a counterexample with 7 points and $k = 3$.

Towards the flip connectivity for special classes of point sets, we consider wheel sets and generalized double circles and prove that the flip graph is connected in both cases:

Theorem 2.11. *Let S be a set of n points in wheel configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter at most $2n - 1$.*

Theorem 2.12. *Let S be a set of n points in generalized double circle configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter $O(n^2)$.*

7.1 A Sufficient Condition

In this section we prove Theorem 7.4 and Theorem 7.5.

Lemma 7.6. *Let S be a point set in general position and $p, q \in S$. Then there exists a path $P \in \mathcal{P}(S)$ which has p and q as its end vertices.*

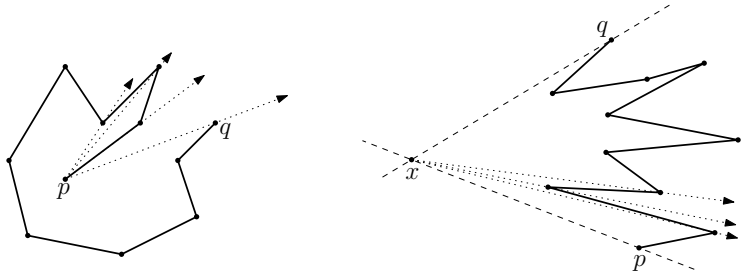

Figure 7.3: For any two given points p and q there exists a plane spanning path having these two points as start and end points. *Left:* the case if at least one of the two points is in the interior of the point set. *Right:* the case when both points are on the boundary of the convex hull.

Proof. We consider two different cases. Let us first assume that at least one of the two points, w.l.o.g. say p, is not an extreme point of S. Sort all other points of S radially around p, starting at q. Connect p to the second point in this order (the point radially just after q) and connect all other points in radial order to a path, such that q becomes the last point of this path; see Figure 7.3 (left). Since p is an interior point, each edge of this path stays in a cone defined by p and two successive points in the radial order, in particular all these cones are disjoint. Hence, we have obtained a plane spanning path starting at p and ending at q.

Assume now for the second case that both given points lie on the boundary of the convex hull of S. Then consider two tangents of $\mathrm{CH}(S)$ through p and q, respectively. By the general position assumption of S we can always perturb these two tangents so that they still go through the two given vertices but are not parallel, and thus cross in a point x outside of the convex hull of S; see Figure 7.3 (right). Sort all points radially around x and connect the points in this order to a path. By construction the points p and q are the first and last point in this sorting, and we thus obtain the required path. □

Theorem 7.4. *Conjecture 7.2 implies Conjecture 7.1.*

Proof. Let S be a point set and $P_s, P_t \in \mathcal{P}(S)$. If P_s and P_t have a common endpoint, we can directly apply Conjecture 7.2 and the statement follows. So assume that P_s has the endpoints v_a and v_b, and P_t has the endpoints v_c and v_d, which are all distinct. By Lemma 7.6 there exists a path P_m having the two endpoints v_a and v_c. By Conjecture 7.2 there is a flip sequence from P_s to P_m with the common endpoint v_a, and again by Conjecture 7.2 there is a further flip sequence from P_m to P_t with the common endpoint v_c. This concludes the proof. □

Towards Theorem 7.5, we first have a closer look at what edges form *viable* starting edges. For a given point set S and points $p, q \in S$, we say that pq

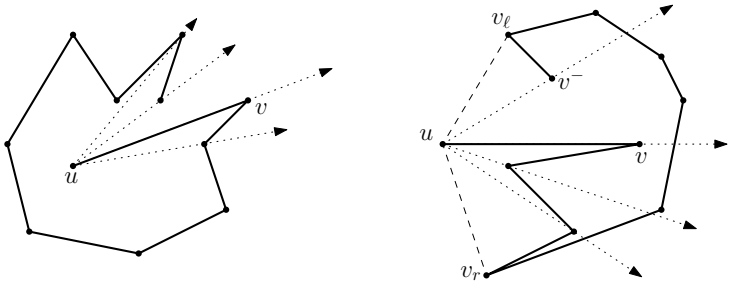

Figure 7.4: Illustration of Lemma 7.7. *Left:* u is an interior point of S. *Right:* u is an extreme point and v an interior point.

forms a *viable starting edge* if there exists a path $P \in \mathcal{P}(S)$ that starts with pq. For instance, an edge connecting two extreme points that are not consecutive along $\mathrm{CH}(S)$ is not a viable starting edge. The following lemma shows that these are the only non-viable starting edges.

Lemma 7.7. *Let S be a point set in general position and $u, v \in S$. The edge uv is a viable starting edge if and only if one of the following is fulfilled:*

(i) u or v lies in the interior of $\mathrm{CH}(S)$, or

(ii) u and v are consecutive along the boundary of $\mathrm{CH}(S)$.

Proof. "⇒". We argue by contraposition. If (i) and (ii) are not fulfilled, i.e., u and v are extreme points that are not consecutive along $\mathrm{CH}(S)$, then there exist points in both halfplanes defined by the line through uv. However, any plane path starting with uv can only reach the points on one side. Hence, uv is not a viable starting edge.

"⇐". If u is in the interior of $\mathrm{CH}(S)$, we sort the remaining points in radial order around u. We construct a path starting with uv that visits the remaining points consecutively in this radial order (see Figure 7.4 (left)).

If u and v are consecutive extreme points, we proceed exactly the same way.

If u is an extreme point and v an interior point, let v_ℓ and v_r be the two neighbors of u along the convex hull. Let $S' \subset S$ be the set of points in the interior of $\mathrm{CH}(S)$ plus $\{v_\ell, v_r\}$. Again we sort the points of S' in radial order around u. Let $v^- \in S'$ be the vertex that is right before v in this radial order.

We construct three paths P_1, P_2, P_3 from v^- via v_ℓ and v_r to v as follows. The paths $P_1 = v_\ell, \ldots, v^-$ and $P_2 = v_r, \ldots, v$ connect the points in S' in radial order (between its corresponding endpoints). Note that P_1 may have length zero. The path $P_3 = v_\ell, \ldots, v_r$ connects v_ℓ to v_r along the boundary of the convex hull of S (along the side not containing u). Then the union of the three paths P_1, P_2, P_3 together with uv forms the desired plane spanning path (see Figure 7.4 (right)). □

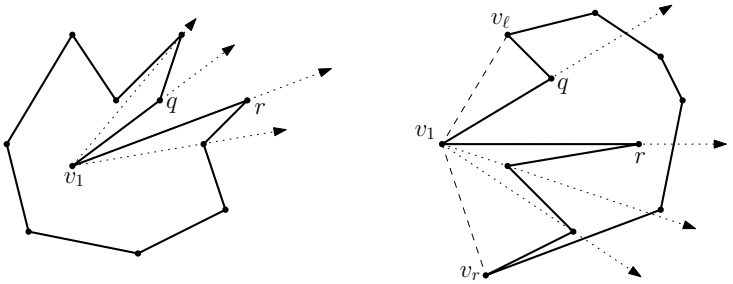

Figure 7.5: There exists a plane spanning cycle of S such that v_1 is connected to two points neighbored in the radial order around v_1. *Left:* v_1 is an interior point. *Right:* v_1 is an extreme point.

The following lemma is the analogue of Lemma 7.6:

Lemma 7.8. *Let S be a point set in general position and $v_1 \in S$. Further let $S' \subset S$ be the set of all points $p \in S$ such that v_1p forms a viable starting edge. Then for two points $q, r \in S'$ that are consecutive in the circular order of S' around v_1, there exists a plane spanning cycle containing the edges v_1q and v_1r.*

Proof. The construction of these plane spanning cycles is completely analogous to the construction of the paths in the proof of Lemma 7.7 but we add the proof for the sake of completeness.

First assume that v_1 is an interior point. Then, by Lemma 7.7, $S' = S \setminus \{v_1\}$ holds. We construct a plane spanning path starting with v_1q and visiting the remaining points in circular order around v_1 such that r is the last in this order. Lastly, connect r to v_1 (see Figure 7.5 (right)).

Now, let v_1 be an extreme point. Again we proceed analogously if q and r are the two neighbors of v_1 along $\mathrm{CH}(S)$. Otherwise, by Lemma 7.7, at least one of the vertices q or r is an interior point. Then we construct the same three paths P_1, P_2, P_3 as in Lemma 7.7 (replacing the roles of v, v^- by q and r). Then the union of P_1, P_2, P_3 together with v_1q and v_1r forms the desired cycle (see Figure 7.5 (right)). □

We are now ready to prove Theorem 7.5:

Theorem 7.5. *Conjecture 7.3 implies Conjecture 7.2.*

Proof. Let S be a point set and $v_1 \in S$. Further let $P, P' \in \mathcal{P}(S, v_1)$. If P and P' have the starting edge in common, then we directly apply Conjecture 7.3 and are done. So let us assume that the starting edge of P is v_1v_2 and the starting edge of P' is v_1v_2'. Clearly $v_2, v_2' \in S'$ holds (where S' is the set of points p such that v_1p forms a viable starting edge). Sort the points in S' in radial order around v_1. Further let $v_x \in S'$ be the next vertex after v_2 in this radial order and C be the plane spanning cycle with edges v_1v_2 and v_1v_x, as guaranteed by Lemma 7.8; see Figure 7.6.

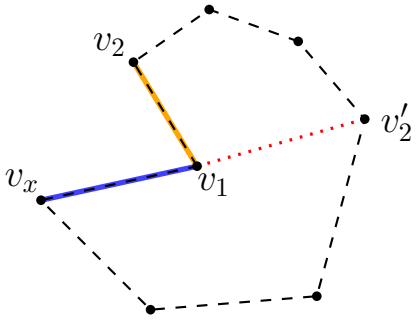

Figure 7.6: Illustratiion of the proof of Theorem 7.5

By Conjecture 7.3, we can flip P to $C \setminus v_1v_x$. Then, flipping v_1v_2 to v_1v_x we get to the path $C \setminus v_1v_2$, which now has v_1v_x as starting edge. We iteratively continue this process of "rotating" the starting edge until reaching v_1v_2'. □

Theorem 7.4 and Theorem 7.5 imply that it suffices to show connectedness of certain subgraphs of the flip graph. A priori it is not clear whether this is an easier or a more difficult task – on the one hand we have smaller graphs, making it easier to handle. On the other hand, we may be more restricted concerning which flips we can perform, or exclude certain "nice" paths.

7.2 Flip Connectivity for Wheel Sets

Akl et al. [AIM07] proved connectedness of the flip graph if the underlying point set S is in convex position. They showed that every path in $\mathcal{P}(S)$ can be flipped to a *canonical path* that uses only edges on the convex hull of S. To generalize this approach to other classes of point sets, we need two ingredients: (i) a set of *canonical paths* that serve as the target of the flip operations and that have the property that any canonical path can be transformed into any other canonical path by a simple sequence of flips, usually of constant length; and (ii) a strategy to flip any given path to some canonical path.

Recall that a set S of $n \geq 4$ points in the plane is a *wheel set* if there is exactly one interior point $c \in S$. We call c the *center* of S and classify the edges on S as follows[2]: an edge incident to the center c is called a *radial edge*, and an edge along $\mathrm{CH}(S)$ is called *spine edge* (the set of spine edges forms the *spine*, which is just the boundary of the convex hull here). All other edges are called *inner edges*. The *canonical paths* are those that consist only of spine edges and one or two radial edges.

We need two observations. Let S be a point set in general position and $P = v_1, \ldots, v_n \in \mathcal{P}(S)$. Further, let v_i $(i \geq 3)$ be a vertex such that no edge on S crosses v_1v_i. We denote the face bounded by $v_1, \ldots, v_i, v_1$ by $\Phi(v_i)$.

[2]Note that we introduced similar notions earlier (boundary edges, diagonal edges, etc.). However, with regard to Section 7.3 it makes sense to introduce a new terminology here.

Observation 7.9. *Let S be a point set in general position, $P = v_1, \dots, v_n \in \mathcal{P}(S)$, and v_i ($i \geq 3$) be a vertex such that no edge on S crosses v_1v_i. Then all vertices after v_i (i.e., $\{v_{i+1}, \dots, v_n\}$) are entirely contained in either the interior or the exterior of $\Phi(v_i)$.*

Observation 7.10. *Let S be a wheel set and let $P = v_1, \dots, v_n \in \mathcal{P}(S)$. Suppose that the edge v_iv_{i+1} of P is an inner edge. Then the sets $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+2}, \dots, v_n\}$ lie on different sides of the line spanned by v_iv_{i+1}.*

Note that the sets $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+2}, \dots, v_n\}$ in Observation 7.10 are non-empty due to Lemma 7.7.

Theorem 2.11. *Let S be a set of n points in wheel configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter at most $2n - 1$.*

Proof. Let $P = v_1, \dots, v_n \in \mathcal{P}(S)$ be a non-canonical path. W.l.o.g., we can assume $v_1 \neq c$ (at least one of the two endpoints of P is not the center). We show how to apply suitable flips to increase the number of spine edges of P.

By Lemma 7.7, the edge v_1v_2 is either radial or a spine edge. We distinguish the two cases:

Case 1 v_1v_2 is radial, i.e., (since we assumed $v_1 \neq c$), we have $v_2 = c$.

Then v_2v_3 is radial and, in analogy to Observation 7.10, the remaining path can only visit vertices on one side of the "line" through $v_1v_2v_3$ (which is bent at v_2); see Figure 7.7(a). Hence, v_3 must be a neighbor of v_1 along the convex hull.

Thus we can increase the number of spine edges by flipping the radial edge v_2v_3 to the spine edge v_1v_3.

Case 2 v_1v_2 is a spine edge.

Let v_a ($a \neq 2$) be the other neighbor of v_1 along the convex hull. Note first that we can assume $v_{a-1}v_a$ to be a spine edge, since otherwise we can increase the number of spine edges by flipping $v_{a-1}v_a$ to v_1v_a. Furthermore, we can assume $a \neq n$, since otherwise we can insert v_1v_n and remove an arbitrary (non-spine) edge (performing a Type 2 flip).

By Observation 7.10, P cannot have any inner edge e before v_a, since otherwise e would have the neighbors v_1 and v_a on the same side. On the other hand, since v_1v_2 and $v_{a-1}v_a$ are spine edges, v_{a+1} must be in the interior of the face $\Phi(v_a)$. Then, by Observation 7.9, all extreme points must already be covered before v_{a+1} (see Figure 7.7(b)). This, however, implies that P does not contain an inner edge.

Hence, if v_1v_2 and $v_{a-1}v_a$ are spine edges P is in fact already canonical.

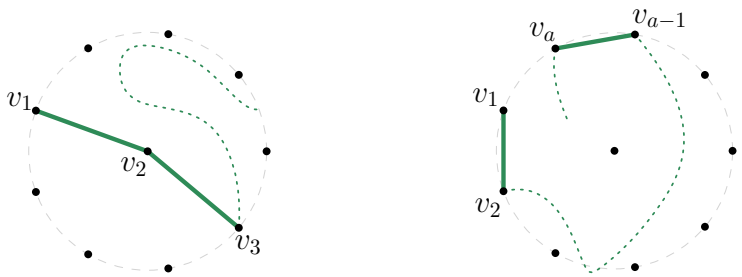

(a) Case 1. **(b)** Case 2.

Figure 7.7: (a) If v_1 and v_3 are not neighbors, there are vertices that cannot be reached. (b) If $v_{a-1}v_a$ is a spine edge, the path after v_a cannot visit any extreme point anymore.

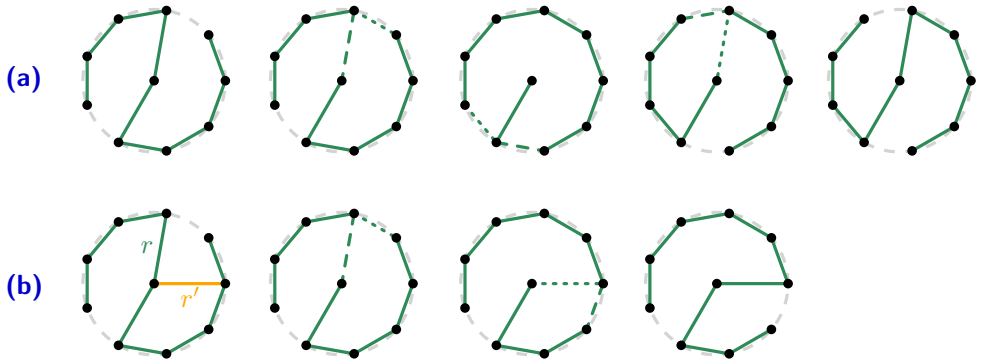

Figure 7.8: (a) Illustration of Remark 1. (b) Illustration of Remark 2.

Iteratively applying above procedure, we obtain a canonical path. We claim that any two canonical paths can be transformed into each other by at most 7 flips. First note that for any fixed set of one or two radial edges there are exactly two different canonical paths in $\mathcal{P}(S)$ using these radial edges.

Remark 1: If two canonical paths C_1, C_2 use the same radial edges, they can be transformed into each other by either one flip (if C_1 and C_2 contain exactly one radial edge) or by three flips (if C_1 and C_2 contain exactly two radial edges.); see Figure 7.8(a).

Remark 2: For every radial edge r' and every canonical path C with radial edge r, using at most two flips, C can be transformed into a canonical path C' such that (i) C' contains r' but not r and (ii) all other radial edges remain the same; see Figure 7.8(b).

Applying Remark 2 at most two times and next Remark 1, it follows that any two canonical paths can be transformed into each other by at most 7 flips.

Lastly, we need to count the number of iterations it takes to transform a plane spanning path into a canonical path. By Observation 7.10 and the

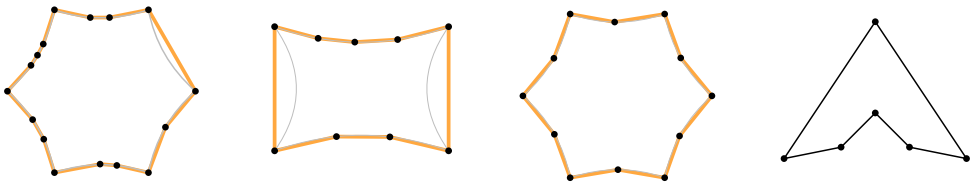

(a) **(b)** Double chain. **(c)** Double circle. **(d)**

Figure 7.9: (a-c) Examples of generalized double circles (the uncrossed spanning cycle is depicted in orange). (d) A point set that is *not* a generalized double circle, but admits an uncrossed spanning cycle.

fact that not all edges can be radial, we know that any plane spanning path contains at least one spine edge and hence, we need at most $n-3$ iterations to reach a canonical path. Since every iteration requires at most 2 flips, we need at most $2(n-3)+7 = 2n+1$ flips in total. However, having a closer look at the intermediate step that transforms the two canonical paths, we see that in the case of $n-3$ iterations where only one radial edge remains, we need only 3 flips rather than 7. Hence, we either save one iteration or 4 intermediate flips, where the former yields a total of $2(n-4)+7 = 2n-1$ flips. □

7.3 Flip Connectivity for Generalized Double Circles

The proof for generalized double circles is generally similar to the one for wheel sets but much more involved. For a point set S and two extreme points $p, q \in S$, we call a subset $CC(p,q) \subset S$ *concave chain* (chain for short) for S, if (i) $p, q \in CC(p,q)$; (ii) $CC(p,q)$ is in convex position; (iii) $CC(p,q)$ contains no other extreme points of S; and (iv) every line ℓ_{xy} through any two points $x, y \in CC(p,q)$ has the property that all points of $S \setminus CC(p,q)$ lie in the open halfplane bounded by ℓ_{xy} that contains neither p nor q. Note that the extreme points p and q must necessarily be consecutive along $\mathrm{CH}(S)$. If there is no danger of confusion, we also refer to the spanning path from p to q along the convex hull of $CC(p,q)$ as the concave chain.

A point set S is in *generalized double circle* position if there exists a family of concave chains such that every inner point of S is contained in exactly one chain and every extreme point of S is contained in exactly two chains. Note that such a family of concave chains is not necessarily unique (e.g. with 4 points in non-convex position, the interior point may belong to any of the three chains). Whenever concerned with a point set in generalized double circle position, we implicitly assume a family of concave chains to be fixed. We denote the class of generalized double circles by GDC (see Figure 7.9 for an illustration). For $S \in \mathrm{GDC}$, it is not hard to see that the union of the concave chains forms an uncrossed spanning cycle:

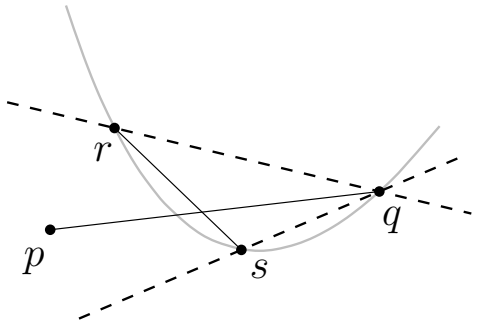

Figure 7.10: Illustration of the proof of Lemma 7.11.

Lemma 7.11. *Every point set $S \in$ GDC admits an uncrossed spanning cycle formed by the union of the concave chains.*

Proof. Let $S \in$ GDC and denote the extreme points of S by $u_1, \ldots, u_k$ in circular order. Assume, for the sake of contradiction, that there is an edge rs along the convex hull of a concave chain $CC(u_i, u_{i+1})$ that is crossed by some edge pq (in particular, at least one of the points r, s is not an extreme point in S). Then p and q lie on different sides of the line ℓ_{rs} through rs and hence, at least one of the points p or q belongs to $CC(u_i, u_{i+1})$, say $q \in CC(u_i, u_{i+1})$. Furthermore, ℓ_{qs} or ℓ_{qr} must have p on the same side as u_i or u_{i+1} (note that this is also the case if r, s, or q coincide with u_i or u_{i+1}); see Figure 7.10. Hence, also $p \in CC(u_i, u_{i+1})$ holds. However, since rs is a hull edge of $CC(u_i, u_{i+1})$ it cannot be crossed by pq; a contradiction. □

Before diving into the details of the proof of Theorem 2.12, we start by collecting preliminary results in a slightly more general setting, namely for point sets S fulfilling the following property:

(USC) There is an *uncrossed* spanning cycle C on S, i.e., no edge joining two points of S crosses any edge of C.

A point set fulfilling (USC) is called *spinal point set*. When considering a spinal point set S, we first fix an uncrossed spanning cycle C, which we call *spine* and all edges in C *spine edges*. For instance, generalized double circles are spinal point sets and the spine is precisely the uncrossed spanning cycle formed by the concave chains as described in Lemma 7.11. Whenever speaking of the spine or spine edges for some point set without further specification, the underlying uncrossed cycle is either clear from the context, or the statement holds for any choice of such a cycle. Furthermore, we call all edges in the exterior/interior of the spine *outer/inner edges*.

Except Lemma 7.13, all preliminary results from Observation 7.12 to Lemma 7.15 hold for the more general setting of spinal point sets.

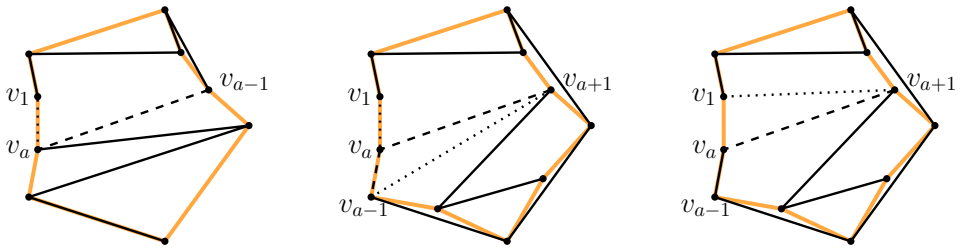

Figure 7.11: Illustration of the three flips in Observation 7.12. The spine is depicted in orange and edge flips are indicated by replacing dashed edges for dotted (in the middle, the two flips are of course executed one after the other).

We define the *canonical paths* to be those that consist only of spine edges. Note that this definition also captures the canonical paths used by Akl et al. [AIM07], and that any canonical path can be transformed into any other by a single flip (of Type 2). Two vertices incident to a common spine edge are called *neighbors*.

Next, we need a strategy to flip an arbitrary path to a canonical path. The biggest issue is to ensure that all intermediate paths in the flip sequence are plane. To this end, it is very helpful if many triangles that are spanned by a spine edge and a third point of S are *empty*, i.e., do not contain any other point from S in their interior.

Valid flips. We collect a few observations which will be useful to confirm the validity of a flip. Whenever we apply more than one flip, the notation in subsequent flips refers to the original graph and not the current (usually we apply one or two flips in a certain step). Figure 7.11 gives an illustration of the following observation.

Observation 7.12. *Let S be a spinal point set, $P = v_1, \dots, v_n \in \mathcal{P}(S)$, and v_1, v_a ($a \neq 2$) be neighbors. Then the following flips are valid (under the specified additional assumptions):*

(a) *flip $v_{a-1}v_a$ to v_1v_a*

(b) *flip in (a) + flip v_av_{a+1} to $v_{a-1}v_{a+1}$* *(if the triangle $\Delta v_{a-1}v_av_{a+1}$ is empty)*

(c) *flip v_av_{a+1} to v_1v_{a+1}* *(if the triangle $\Delta v_1v_av_{a+1}$ is empty and $v_{a-1}v_a$ is a spine edge).*

Strictly speaking, in Observation 7.12(c) we do not require $v_{a-1}v_a$ to be a spine edge, but merely to be an edge not crossing v_1v_{a+1}. The following lemma provides structural properties for *generalized double circles*, if the triangles in Observation 7.12(b,c) are non-empty:

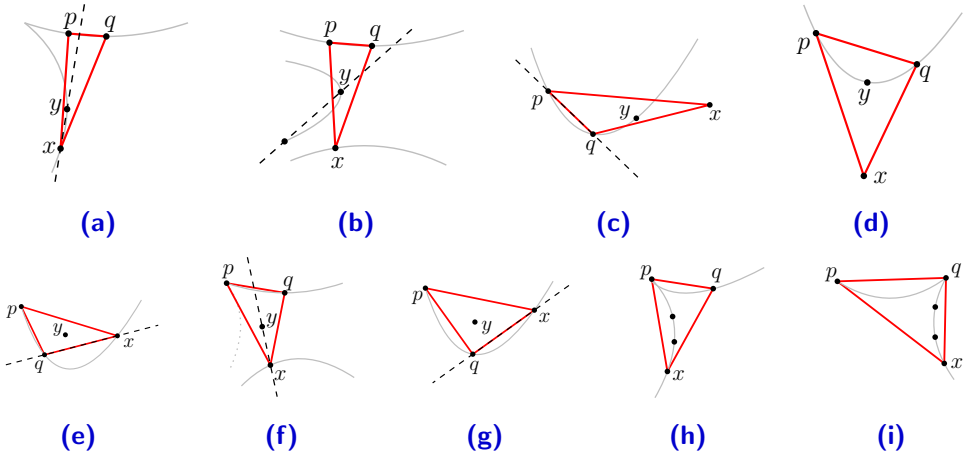

Figure 7.12: Illustration of the proof of Lemma 7.13.

Lemma 7.13. *Let $S \in$ GDC and $p, q, x \in S$ such that p and q are neighbors. Further, let the triangle Δpqx be non-empty. Then the following holds:*

(i) At least one of the two points p, q is an extreme point (say p),

(ii) x does not lie on a common chain with p and q, but shares a common chain with either p or q (the latter may only happen if q is also an extreme point).

Proof. Concerning part (i), assume for the sake of contradiction that neither p nor q is an extreme point and let $y \in \Delta pqx$. First recall that p and q belong to exactly one concave chain and since they are neighbors, they belong to the same chain – call it C_{pq}. Further note that y is not an extreme point and hence, also belongs to a unique concave chain.

Let us first consider the case that x and y share a common chain C_{xy}. Since concave chains are in convex position, this cannot be C_{pq}. However, since y is contained in the interior of Δpqx, the line ℓ_{xy} through x and y separates p and q; a contradiction to property (iv) of the chain C_{xy} (see Figure 7.12(a)).

Consider next, that x and y do not share a common chain. If y does not belong to C_{pq}, then neither of the points p, q, x belongs to the chain C_y of y. However, any line through y and another point of C_y separates these three points; a contradiction to property (iv) of C_y (see Figure 7.12(b)). If y belongs to C_{pq}, then either p and q are not neighbors or the line ℓ_{pq} has x and the extreme points (in S) of C_{pq} on the same side (see Figure 7.12(c) and Figure 7.12(d)); a contradiction either way. This finishes the proof of part (i).

Concerning part (ii), first note that the extreme point p belongs to two chains C_{pq} and $C_{p\bar{q}}$ (the former also contains q while the latter does not). Also assume that q is not an extreme point. Then we need to show that x belongs to $C_{p\bar{q}}$.

First we argue that x does not belong to C_{pq}. Indeed, if this was the case, the point $y \in \Delta pqx$ (which does not belong to C_{pq}) would lie on the same side of ℓ_{qx} as p (see Figure 7.12(e)); a contradiction to property (iv) of C_{pq}.

Second, suppose for the sake of contradiction that x belongs to some other chain $C \neq C_{p\bar{q}}$. Again we consider the two cases where x and y share a common chain or not.

First, let x and y share a common chain (which, as before, is not C_{pq}). Since ℓ_{xy} separates p and q this is a contradiction, except p, x, y lie on a common chain, which however is excluded because $x \notin C_{p\bar{q}}$ (see Figure 7.12(f)).

Second, let x and y not share a common chain. The cases that y belongs to a chain that is not $C_{p\bar{q}}$ are completely analogous as before (see Figures 7.12(b,c,d)). Hence, it remains to consider $y \in C_{p\bar{q}}$. In this case, the line ℓ_{py} separates q and x; again a contradiction (see Figure 7.12(g)).

The case that q is an extreme point is now completely analogous. Figures 7.12(h,i) summarize the possible configurations for a non-empty triangle Δpqx. □

Combinatorial distance measure. In contrast to the proof for wheel sets, it may now not be possible anymore to directly increase the number of spine edges and hence, we need a more sophisticated measure. Let C be the spine of a spinal point set S and $p, q \in S$. Further let $o \in \{\text{cw}, \text{ccw}\}$ be an orientation. We define the *distance* between p, q *in direction* o, denoted by $d^o(p, q)$, as the number of spine edges along C that lie between p and q in direction o. Furthermore, we define the *distance* between p and q to be

$$d(p, q) = \min\{d^{\text{cw}}(p, q), d^{\text{ccw}}(p, q)\}.$$

Note that neighboring points along the spine have distance one. Using this notion we define the *weight* of an edge to be the distance between its endpoints and the (overall) weight of a graph on S to be the sum of its edge weights.

Our goal is to perform weight decreasing flips (the terms increase/decrease always refer to *strict* increase/decrease). To this end, we state two useful preliminary results (see also Figure 7.13):

Observation 7.14. *Let S be a spinal point set and p, q, r be three neighboring points in this order (i.e., q lies between p and r). Further let $s \in S \setminus \{p, q, r\}$ be another point. Then $d(p, s) < d(q, s)$ or $d(r, s) < d(q, s)$ holds.*

Combining Observation 7.12 and Observation 7.14, it is apparent that we can perform weight decreasing flips whenever $\Delta v_{a-1} v_a v_{a+1}$ and $\Delta v_1 v_a v_{a+1}$ are empty.

Lemma 7.15. *Let S be a spinal point set, $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ a non-canonical path, and v_a, v_b be the neighbors of v_1 as well as v_c, v_d be the neighbors of v_n. If $\max(a, b) > \min(c, d)$, then the number of spine edges in P can be increased by performing at most two flips, which also decreases the overall weight of P.*

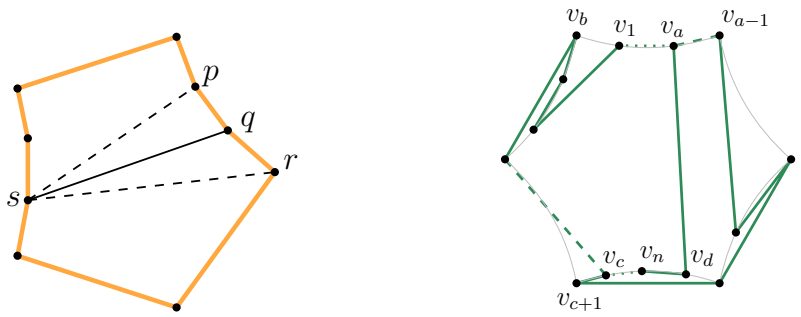

(a) Observation 7.14. **(b)** Lemma 7.15.

Figure 7.13: (a) One of the dashed edges has smaller weight than the solid: $d(s,q) = 4$; $d(s,p) = 4$; $d(s,r) = 3$. (b) The initial path is depicted by solid and dashed edges. Flipping the dashed edges to the dotted edges increases the number of spine edges.

Proof. First of all we can assume v_1 and v_n not to be neighbors, since otherwise we can increase the number of spine edges by a single flip, inserting v_1v_n and removing an arbitrary non-spine edge. Let $\max(a,b) > \min(c,d)$ and assume w.l.o.g. that $a > c$ holds. In particular $a \neq 2$ and $c \neq n-1$ holds.

If $v_{a-1}v_a$ or v_cv_{c+1} is not a spine edge, we can increase the number of spine edges by a single flip using Observation 7.12(a). Hence, we can assume these edges to be spine edges, which implies that v_av_{a+1} and $v_{c-1}v_c$ are not spine edges.

Then we flip the spine edge $v_{a-1}v_a$ to the other spine edge v_1v_a resulting in the path $v_{a-1},\dots,v_{c+1},v_c,\dots,v_1,v_a,\dots,v_n$. Next, flipping $v_{c-1}v_c$ to v_cv_n is valid and now replaces a non-spine edge by a spine edge. See Figure 7.13(b) for an illustration. □

Lemma 7.15 essentially enables us to perform weight decreasing flips whenever the path traverses a neighbor of v_n before it reached both neighbors of v_1. We are now ready to prove Theorem 2.12, but briefly summarize the proof strategy from a high-level perspective beforehand:

High-level proof strategy. To flip an arbitrary path $P \in \mathcal{P}(S)$ to a canonical path, we perform iterations of suitable flips such that in each iteration we either

(i) strictly increase the number of spine edges along P, while not increasing the overall weight of P, or

(ii) strictly decrease the overall weight of P, while not decreasing the number of spine edges along P.

Note that for the connectivity of the flip graph it is not necessary to guarantee the non increasing overall weight in the first part. However, this will provide us with a better bound on the diameter of the flip graph. The remainder of this chapter is concerned with the proof of Theorem 2.12.

Theorem 2.12. *Let S be a set of n points in generalized double circle configuration. Then the flip graph on $\mathcal{P}(S)$ is connected with diameter $O(n^2)$.*

Proof. Let $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ be a non-canonical path. As described in the high-level proof strategy, we show how to iteratively transform P to a canonical path by increasing the number of spine edges or decreasing its overall weight. As usual let v_a ($a \neq 2$) be a neighbor of v_1.

We can assume, w.l.o.g., that v_1 and v_n are not neighbors, since otherwise we can flip an arbitrary (non-spine) edge of P to the spine edge v_1v_n (performing a Type 2 flip), i.e., $a < n$. Furthermore, we can also assume w.l.o.g., that $v_{a-1}v_a$ is a spine edge, since otherwise we can flip $v_{a-1}v_a$ to the spine edge v_1v_a (Observation 7.12(a)). This also implies that the edge v_av_{a+1}, which exists because $a < n$, is not a spine edge, since v_a already has the two neighbors v_{a-1} and v_1.

We distinguish two cases – v_1v_2 being a spine edge or not:

Case 1 v_1v_2 is not a spine edge.

This case is easier to handle, since we are guaranteed that both neighbors v_a, v_b ($a, b \neq 2$) of v_1 are potential candidates to flip to. In order to apply Observation 7.12, we require $\Delta v_1v_av_{a+1}$ or $\Delta v_1v_bv_{b+1}$ to be empty. So, let's consider these two subcases separately:

Case 1.1 $\Delta v_1v_av_{a+1}$ (or analogously $\Delta v_1v_bv_{b+1}$) is empty.

Then we apply the following two flips:

$$\text{flip } v_av_{a+1} \text{ to } v_1v_{a+1} \qquad \text{and} \qquad \text{flip } v_1v_2 \text{ to } v_1v_a,$$

where the first flip results in the path $v_a, \ldots, v_1, v_{a+1}, \ldots, v_n$ (and is valid by Observation 7.12(c)) and the second flip then results in the path $v_2, \ldots, v_a, v_1, v_{a+1}, \ldots, v_n$ (valid due to Observation 7.12(a)). The first flip replaces a non-spine edge by another non-spine edge and may increase the weight by at most one. The second flip replaces a non-spine by a spine edge, which also decreases the weight by at least one. Together, the number of spine edges increases, while the overall weight does not increase.

Case 1.2 $\Delta v_1v_av_{a+1}$ and $\Delta v_1v_bv_{b+1}$ are not empty.

Lemma 7.13(i) implies that (v_1 or v_a) and (v_1 or v_b) is an extreme point. We consider these cases separately. For convenience we introduce one more term: For an outer edge e (which necessarily

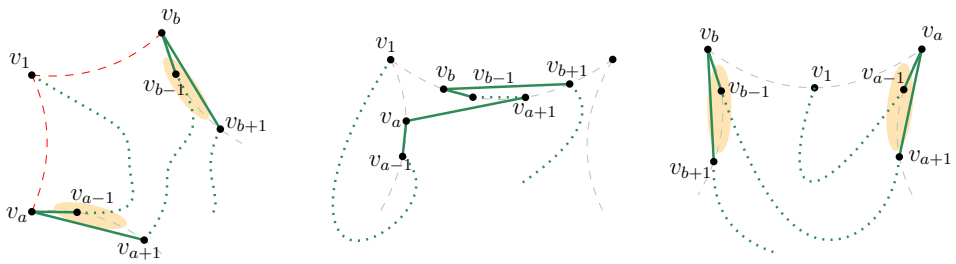

(a) Case 1.2.1. **(b)** Case 1.2.3. **(c)** Case 1.2.4.

Figure 7.14: (a) Case 1.2.1. There cannot be more than $n/2$ points in each yellow region and hence, at least one of the flips $v_a v_{a+1}$ to $v_{a-1}v_{a+1}$ or $v_b v_{b+1}$ to $v_{b-1}v_{b+1}$ decreases the weight. (b) Case 1.2.3. Either $\Delta v_1 v_a v_{a+1}$ or $\Delta v_1 v_b v_{b+1}$ must be empty. (c) Case 1.2.4.

connects two points from the same chain), we say that the points along this chain between the two endpoints of e (excluding the endpoints) lie *under* e.

Case 1.2.1 v_1, v_a, v_b are all extreme points.
The triangles $\Delta v_1 v_a v_{a+1}$ and $\Delta v_1 v_b v_{b+1}$ being non-empty implies that both, $v_a v_{a+1}$ and $v_b v_{b+1}$ are outer edges (see Figure 7.14(a)). Not both of them can contain $n/2$ points under its edge. Hence, one of the two flips (replacing $v_a v_{a+1}$ by $v_{a-1}v_{a+1}$, or $v_b v_{b+1}$ by $v_{b-1}v_{b+1}$) is weight decreasing, say w.l.o.g. the one for v_a. Then we apply the following flips:

$$\text{flip } v_{a-1}v_a \text{ to } v_1 v_a \qquad \text{and} \qquad \text{flip } v_a v_{a+1} \text{ to } v_{a-1}v_{a+1}.$$

The first flip is valid due to Observation 7.12(a) (replacing a spine edge by another spine edge) and the second is valid due to Observation 7.12(b) and decreases the weight by assumption.

Case 1.2.2 v_1 and v_a are extreme points and v_b is not (analogous with exchanged roles of v_a and v_b).
In this case, using Lemma 7.13(ii), we conclude that the triangle $\Delta v_1 v_b v_{b+1}$ is empty; a contradiction.

Case 1.2.3 v_1 is an extreme point and v_a, v_b are not.
By Lemma 7.13(ii), the edge $v_a v_{a+1}$ must be an inner edge between the two concave chains of v_1. If v_{a+1} is a neighbor of v_1, we can simply replace $v_a v_{a+1}$ by $v_1 v_{a+1}$ (Observation 7.12(a)). Otherwise, the other neighbor v_b of v_1 cannot be incident to an inner edge and hence, $\Delta v_1 v_b v_{b+1}$ is empty; a contradiction (see Figure 7.14(b)).

Case 1.2.4 v_1 is not an extreme point and v_a, v_b are both extreme points.

This case is analogous to Case 1.2.1. Again, both, $v_a v_{a+1}$ and $v_b v_{b+1}$ must be outer edges and the regions under those two edges must be disjoint (see Figure 7.14(c)). Hence, one of the two flips (replacing $v_a v_{a+1}$ by $v_{a-1} v_{a+1}$, or $v_b v_{b+1}$ by $v_{b-1} v_{b+1}$) is weight decreasing, say w.l.o.g. the one for v_a. Then, as in Case 1.2.1, we apply the following flips:

$$\text{flip } v_{a-1}v_a \text{ to } v_1 v_a \qquad \text{and} \qquad \text{flip } v_a v_{a+1} \text{ to } v_{a-1} v_{a+1}.$$

Case 2 $v_1 v_2$ is a spine edge.

In this case we will consider P from both its endpoints v_1 and v_n. Our general strategy here is to first rule out some easier cases and collect all those cases where we cannot immediately make progress. For these remaining "bad" cases we consider the setting from both ends of the path, i.e., we need to consider all combinations of bad cases.

Case 2.1 v_1, v_a are not extreme points. And v_{a-1} is also not an extreme point. Then we either

- flip $v_a v_{a+1}$ to $v_1 v_{a+1}$, or
- flip $v_{a-1} v_a$ to $v_1 v_a$ and $v_a v_{a+1}$ to $v_{a-1} v_{a+1}$.

By Observation 7.14, one of the two choices decreases the overall weight and all flips are valid because neither of the vertices v_1, v_a, v_{a-1} is an extreme point and hence, the triangles $\Delta v_1 v_a v_{a+1}$ and $\Delta v_{a-1} v_a v_{a+1}$ are empty (Lemma 7.13).

Case 2.2 v_1, v_a are not extreme points. And v_{a-1} is an extreme point.

Case 2.2.1 v_{a-2} lies on the same chain as v_1.

Case 2.2.1.1 $d(v_{a-2} v_{a-1}) = 2$.

This implies v_{a-2} to be an extreme point and there is only one other point in S that does not belong to the chain of v_{a-2} and v_{a-1}; see Figure 7.15(a). This, however, implies that both triangles $\Delta v_1 v_a v_{a+1}$ and $\Delta v_{a-1} v_a v_{a+1}$ are empty and hence, using Observation 7.12 and Observation 7.14, we can perform weight decreasing flips.

Case 2.2.1.2 $d(v_{a-2} v_{a-1}) > 2$.

Then we flip $v_{a-2} v_{a-1}$ to $v_1 v_{a-1}$, which decreases the overall weight, since $v_1 v_{a-1}$ has weight 2. See Figure 7.15(b).

Case 2.2.2 v_{a-2} lies on a different chain as v_1.

Recall that the notation $\Phi(v_x)$ refers to the face bounded by $v_1, \ldots, v_x, v_1$.

If v_{a+1} lies in the exterior of $\Phi(v_a)$, then $v_{a-1}, v_a, v_1, v_{a+1}$ share a common chain and we can apply the flips of Case 2.1.

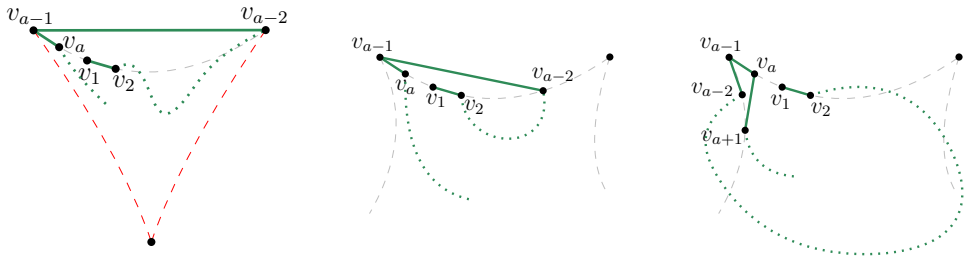

(a) Case 2.2.1.1. **(b)** Case 2.2.1.2. **(c)** Case 2.2.2.

Figure 7.15: (b) Replacing $v_{a-2}v_{a-1}$ by v_1v_{a-1} is weight decreasing. (c) The first "bad" case.

If v_{a+1} lies in the interior of $\Phi(v_a)$ this constitutes our first "bad" case (note that Observation 7.9 implies that the subpath $v_{a+1}, \ldots, v_n$ does not contain any extreme point). Further, we can assume $v_{a-2}v_{a-1}$ also to be a spine edge, since otherwise we flip $v_{a-1}v_a$ to v_1v_a and are back in Case 1. Also we can assume v_av_{a+1} to be an inner edge towards the v_{a-1} chain, since otherwise we can make progress (using Observations 7.12 and 7.14) because both triangles $\Delta v_1v_av_{a+1}$ and $\Delta v_{a-1}v_av_{a+1}$ would be empty by Lemma 7.13(ii). See Figure 7.15(c).

Case 2.3 v_1 is an extreme point and v_a is not an extreme point.

Again it suffices to consider the case that at least one of the triangles $\Delta v_1v_av_{a+1}$ or $\Delta v_{a-1}v_av_{a+1}$ is non-empty. Either way, again using Lemma 7.13(ii), this can only happen, if v_{a+1} is in the interior of $\Phi(v_a)$, more precisely, v_av_{a+1} is an inner edge towards one of the neighboring chains of the v_a chain (where the one case is only relevant if v_{a-1} is an extreme point); see Figure 7.16(a,b). This constitutes the second "bad" case. As before, there are no extreme points along the subpath $v_{a+1}, \ldots, v_n$.

Case 2.4 v_a is an extreme point and v_1 is not an extreme point.

If v_av_{a+1} is an inner edge both triangles $\Delta v_1v_av_{a+1}$ and $\Delta v_{a-1}v_av_{a+1}$ are empty and we can make progress by Observation 7.12 and Observation 7.14. Hence, there are two "bad" cases to consider here, namely when v_av_{a+1} is an outer edge (see Figure 7.16(c,d)).

Case 2.5 v_1, v_a are extreme points.

Again, there are two "bad" cases to consider here, namely when v_av_{a+1} is an outer edge or an inner edge to the v_1, v_2 chain (see Figure 7.16(e,f)).

Let us summarize the six "bad" cases, where we cannot immediately make progress (note that (II) comprises two cases but for the following arguments it will not be important to distinguish between them):

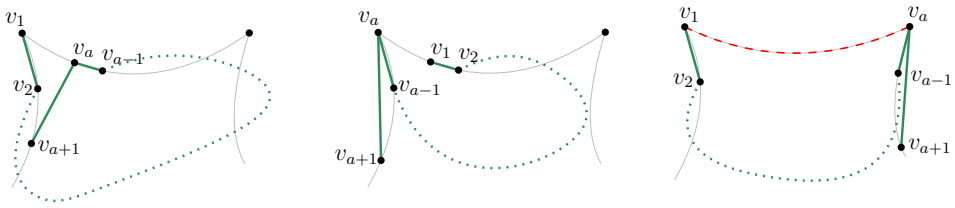

(a) Case 2.3. **(c)** Case 2.4. **(e)** Case 2.5.

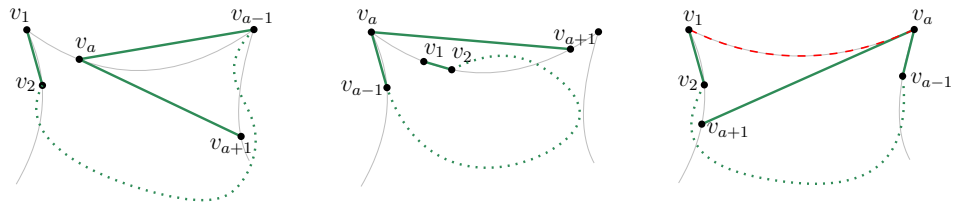

(b) Case 2.3. **(d)** Case 2.4. **(f)** Case 2.5.

Figure 7.16: (a,b) The two "bad" cases of Case 2.3. Either way, $v_a v_{a+1}$ is in the interior of $\Phi(v_a)$. (c,d) The two "bad" cases of Case 2.4. The edge $v_a v_{a+1}$ is in the exterior of $\Phi(v_a)$. (e,f) The two "bad" cases of Case 2.5. The edge $v_a v_{a+1}$ can be in the interior (f) or exterior (e) of $\Phi(v_a)$. Recall that the dashed red arc shall emphasize that there is no vertex lying on this chain between the two extreme points v_1, v_a.

(I) v_1, v_a are not extreme points, v_{a-1} is an extreme point, $v_a v_{a+1}$ is an inner edge towards the v_{a-1} chain that does not contain v_1, and $v_{a-1}v_{a-2}$ is a spine edge (Figure 7.15(c)).

(II) v_1 is an extreme point, v_a is not an extreme point, and $v_a v_{a+1}$ is an inner edge towards a neighboring chain (in the one case v_{a-1} must be extreme); Figure 7.16(a,b).

(IIIa) v_1 is not an extreme point, v_a is an extreme point, and $v_a v_{a+1}$ is an outer edge on the v_{a-1} chain (Figure 7.16(c)).

(IIIb) v_1 is not an extreme point, v_a is an extreme point, and $v_a v_{a+1}$ is an outer edge on the v_1 chain (Figure 7.16(d)).

(IVa) v_1, v_a are extreme points and $v_a v_{a+1}$ is an outer edge (Figure 7.16(e)).

(IVb) v_1, v_a are extreme points and $v_a v_{a+1}$ is an inner edge to the v_1 chain (Figure 7.16(f)).

In the remainder of the proof we settle these "bad" cases by arguing about both ends of the path, i.e, we consider all $\binom{6}{2} + 6 = 21$ combinations of these

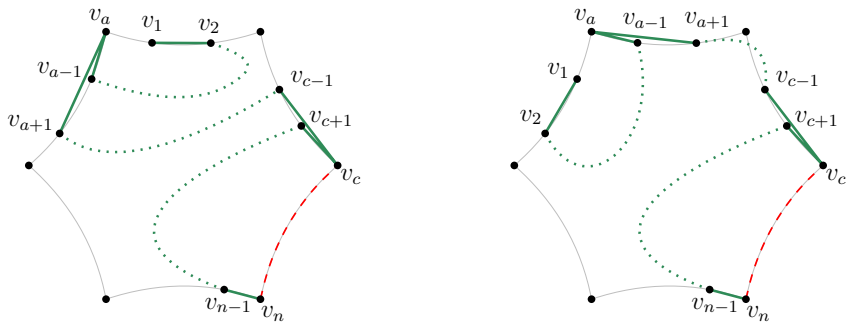

Figure 7.17: The two non-isomorphic ways to realize a combination of "bad" cases ((IIIa)+(IVa) is drawn here). Recall that the red dashed arc indicates that there are no more vertices on this chain.

"bad" cases. Recall that the order at the v_n end must be thought of as inverted, e.g., the analogue of $v_{a-1}v_a$ is v_cv_{c+1} and the analogue of v_av_{a+1} is $v_{c-1}v_c$.

Further note that there are two non-isomorphic ways two realize a combination of "bad" cases (see Figure 7.17). Furthermore, there are two ways to connect the "loose" ends of the fixed structures with each other (one case where we have $a < c$ and another where we have $c < a$). However, by Lemma 7.15, only the former is of interest.

Since we are not using the edges of the intermediate subpaths, we are mostly interested in the question whether or not a certain combination can be realized in a plane manner (in one of the two ways) or not.

Fortunately, we can immediately exclude several combinations with the following observations:

1. As noted above, we can assume $a < c$ and hence, no "bad" case where v_{a+1} is in the interior of $\Phi(v_a)$ can be combined with a "bad" case having v_n or v_c as extreme point. Hence, we can exclude all combinations involving (I), (II), or (IVb), except (I)+(I). In fact, (I)+(I) is also not possible as the path starting from v_1 would first need to traverse v_{c+1}, then v_a and then v_c, which is not possible.

2. If v_av_{a+1} is an outer edge, there is always an "easy" flip possible to either v_1 or v_{a-1} (depending on which lies on the same chain as v_a and v_{a+1}). Hence, similar to Case 1.2.4, we can rule out cases by showing that the regions (towards we flip v_av_{a+1} and $v_{c-1}v_c$ to) are disjoint. Also, we need to consider a subtle edge case, namely when v_av_{a+1} and $v_{c-1}v_c$ coincide. We consider the remaining six cases separately (see also Figure 7.18). As depicted in Figure 7.18, we have disjoint regions in all 6 remaining combinations except for the edge cases where the edges v_av_{a+1} and $v_{c-1}v_c$ coincide.

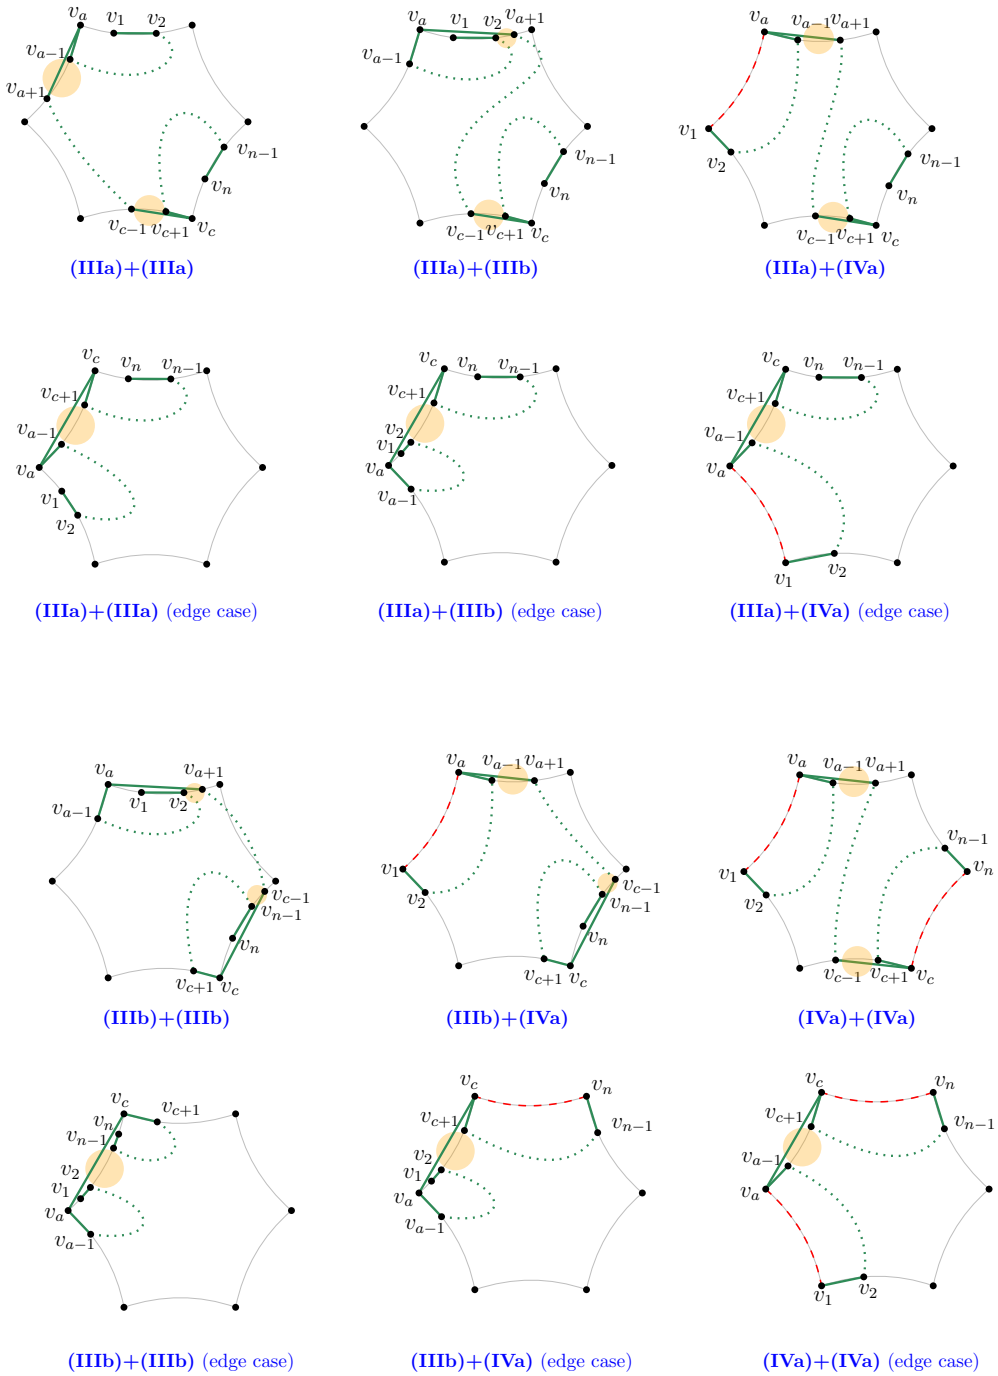

Figure 7.18: The remaining combinations where we try to find disjoint regions each having to contain at least half of the point set. Edge cases are those cases where the edges $v_a v_{a+1}$ and $v_{c-1} v_c$ coincide.

Hence, 1. and 2. rule out all possible cases, except the 6 edge cases from 2., where also these regions (containing at least $n/2$ points) coincide on the chain $CC(v_a, v_c)$. These can be resolved as follows. In all cases there exists an inner edge $v_{x-1}v_x$ along the path $v_2 \dots v_{a-1}$ that is "closest" to v_1 and v_x lies on the chain $CC(v_a, v_c)$. By closest we mean that v_1v_x does not intersect the path. Hence, flipping $v_{x-1}v_x$ to v_1v_x forms a valid flip and at least one of the two choices (either from the v_1 end or v_n) decreases the weight. This argument is analogous in all six edge cases, with the only difference that we may replace v_1 by v_{a-1} by first flipping $v_{a-1}v_a$ to v_1v_a if v_1 lies on the $CC(v_a, v_c)$ chain. This concludes the case distinction. In each case we showed how to perform suitable flips to increase the number of spine edges or decrease the overall weight.

Now, iteratively applying the above process, we will eventually transform P to a canonical path that consists only of spine edges (the only paths with minimum overall weight). Doing the same for Q and noting that any pair of canonical paths can be transformed into each other by a single flip, the connectedness of the flip graph follows.

Concerning the required number of flips, note that any edge has weight at most $\frac{n}{2} - 1$ and the path has $n - 1$ edges. Hence, the total number of iterations to transform P into a canonical path is at most

$$\left((n-1) \cdot \left(\frac{n}{2} - 2\right) + (n-1)\right) \in O(n^2)$$

Furthermore, any iteration requires at most two flips and hence, the total number of flips to transform P into Q is still in $O(n^2)$. □

Chapter 8

Compatible Trees

In the previous chapter we studied the flip graph on *straight-line* paths, whereas in this chapter, we are now leaving the realm of straight-line drawings. For a fixed integer n, let D be a simple drawing of the complete graph K_n and let $\mathcal{T}_D$ be the set of all subdrawings of D which are plane spanning trees. Note that $\mathcal{T}_D$ is non-empty, as it contains at least the n stars in D. In this chapter, unless stated otherwise, the word *tree* always refers to a plane spanning tree in $\mathcal{T}_D$, where the drawing D is either clear from the context or the statement holds for any simple drawing of K_n.

Recall that two trees are *compatible* if their union is plane. Again, we are concerned with the question of connectedness of a certain flip graph. In this chapter, our ground set is $\mathcal{T}_D$ (for some fixed drawing D of K_n) and the flip operation is defined via compatibility, i.e., two vertices are adjacent if and only if the corresponding trees are compatible.

For clarity and easier distinguishability we call this particular flip graph *compatibility graph*.

Question 8.1. *Is the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ connected for every simple drawing D of the complete graph K_n, for all $n \in \mathbb{N}$?*

Note that the compatibility graph is not connected if the underlying drawing is not a drawing of the complete graph but instead one of a complete bipartite graph for example (see Figure 8.1).

Related work. Considering the notion of compatibility, most of the work has been done in the straight-line setting, e.g., in the context of perfect matchings: Houle et al. [HHNR05] proved connectedness of the compatibility graph (even using a weaker notion of flips). Their bound of $n-2$ on the diameter has been improved to $O(\log n)$ by Aichholzer et al. [ABD+09], while Razen [Raz08] provided a lower bound of $\Omega(\log n / \log\log n)$. Aloupis et al. [ABLS15] showed connectedness of the compatibility graph of perfect matchings, where the underlying points are colored red and blue and Aichholzer et al. [ABH+18] provided an upper bound of $2n$ on the diameter in this setting.

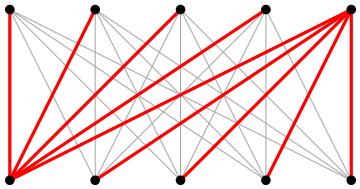

Figure 8.1: The red tree intersects all other edges and hence, is an isolated vertex in the corresponding compatibility graph.

In the context of plane spanning trees, Aichholzer et al. [AAH02] showed, in the straight-line setting, that the compatibility graph is connected with diameter $O(\log n)$ and slightly later the refined upper bound of $O(\log k)$ [AAHK06], where k denotes the number of convex layers of the point set. Buchin et al. [BRUW09] provided a construction proving a lower bound of $\Omega(\log n / \log \log n)$ and a matching lower bound of $\Omega(\log k)$ in terms of the number of convex layers.

It is natural to extend this question, concerning the connectedness of the compatibility graph, also to simple drawings. In this chapter, we aim to shed some light on this wide open topic in the context of plane spanning trees. Note that every plane spanning tree (in every simple drawing of K_n) has a compatible tree, which follows from the property that every maximal plane subgraph is 2-connected [GTP21]: given a plane spanning tree, which clearly is not 2-connected, there must be an edge that can be added and preserves planeness and hence, gives rise to a compatible tree.

Contribution. We approach Question 8.1 from two directions, proving a positive answer for special classes of drawings (namely, cylindrical, monotone, and strongly c-monotone drawings) and for special classes of spanning trees (namely stars, double stars, and twin stars).

Theorem 2.13. *Let D be a cylindrical, monotone, or strongly c-monotone drawing of the complete graph K_n. Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected.*

Theorem 2.14. *Let D be a simple drawing of the complete graph K_n and let $\mathcal{T}_D^*$ be the set of all plane spanning stars, double stars, and twin stars on D. Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D^*)$ is connected.*

Section 8.1 is concerned with the proof of Theorem 2.13, while Section 8.2 is dedicated to the proof of Theorem 2.14.

8.1 Special simple drawings of K_n

In this section we prove connectedness of the compatibility graph for certain classes of drawings. Clearly, for any drawing of K_n that has a plane spanning

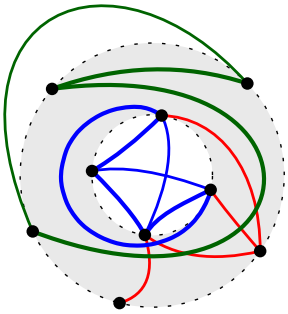

Figure 8.2: Illustration of a cylindrical drawing. Side edges are red; inner edges are blue; and outer edges are green. Cycle edges are drawn thicker.

tree which is not crossed by any edge of D, the compatibility graph is connected with diameter at most 2. This is, for example, the case for *2-page book drawings*, where the vertices are placed along a line and each edge lies entirely in one of the two open halfplanes defined by this line.

8.1.1 Cylindrical drawings

Recall that in a *cylindrical drawing* of a graph the vertices are placed along two concentric circles, the *inner* and *outer* circle, and no edge is allowed to cross these circles. We denote the vertices on the inner/outer circle of a cylindrical drawing by *inner/outer vertices* and the edges connecting any two inner/outer vertices by *inner/outer edges*. The remaining edges are called *side edges*. Furthermore, edges that join consecutive vertices of either of the two circles are called *cycle edges*. See Figure 8.2 for an illustration.

Remark 8.2. *Cycle edges may only be crossed by side edges and both circles each contain at most one such cycle edge that is crossed.*

In other words, there exists an *inner* and an *outer* Hamiltonian path visiting all inner/outer vertices and consisting entirely of cycle edges that are not crossed by any other edge.

Furthermore, any plane spanning tree in a cylindrical drawing contains at least one side edge.

Lemma 8.3. *Let D be a cylindrical drawing of K_n. Then the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter at most 4.*

Proof. Let $T_1, T_2 \in \mathcal{T}_D$ and $e_1 = v_s v_r$ be a side edge of T_1 and $e_2 = v'_s v'_r$ be a side edge of T_2. Furthermore, as guaranteed by Remark 8.2, let $S_1 \in \mathcal{T}_D$ be a plane spanning tree consisting of the two uncrossed paths of cycle edges and e_1 (and similarly $S_2 \in \mathcal{T}_D$ consists of the cycle paths and e_2). Clearly, T_1 and S_1 are compatible as well as T_2 and S_2. If S_1 and S_2 are compatible, we are done. Otherwise, e_1 and e_2 cross in D. As the 4-tuple $\{v_s, v'_s, v_r, v'_r\}$ induces at

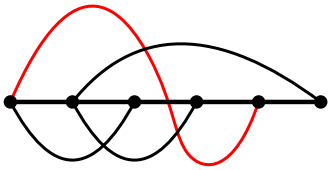

Figure 8.3: Illustration of a monotone drawing. The spine path is drawn thicker. The red edge is a twiggly edge.

most one crossing (due to properties of simple drawings), the two side edges $v_s v'_r$ and $v'_s v_r$ can neither cross e_1 nor e_2. Take one of these edges and call it e' and construct a plane spanning tree S_3, using only e' and the uncrossed cycle paths. Then S_3 is compatible to S_1 as well as S_2 and hence, T_1,S_1,S_3,S_2,T_2 is a path of length 4 in $\mathcal{F}(\mathcal{T}_D)$. □

8.1.2 Monotone drawings

A simple drawing in which no two vertices have the same x-coordinate and every edge is drawn as an x-monotone curve is called *monotone*. Let $v_1, v_2, \ldots, v_n$ denote the vertices in increasing x-order. Then, $\mathcal{S} = v_1, v_2, \ldots, v_n$ is a plane spanning path, which we call *spine path*. An edge that intersects the spine path is called *twiggly* edge. See Figure 8.3 for an illustration.

We define a relation on the twiggly edges of D as follows: for two twiggly edges e, f we have $e \succ f$ if they are non-intersecting and there is a vertical line intersecting the relative interiors of both edges that intersects e at a larger y-coordinate than f. All other pairs of twiggly edges are incomparable, i.e., those that have disjoint support or are intersecting. Note that the relation $\succ$ does not define a poset, since transitivity is not fulfilled.

For a set E of pairwise non-intersecting twiggly edges, an edge $e \in E$ is *maximal* in E, if there is no other edge $f \in E$ with $f \succ e$. Note that this relation is acyclic, i.e., there are no twiggly edges $e_1, \ldots, e_k$ such that $e_1 \succ e_2 \succ \ldots \succ e_k \succ e_1$. Hence, any non-empty set of twiggly edges admits a maximal element.

Lemma 8.4. *For any monotone drawing D of K_n, the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter $O(n)$.*

Proof. We prove the connectedness by showing that any plane spanning tree in D can be transformed (via a sequence of compatible trees) to the spine path $\mathcal{S}$. So, let $T \in \mathcal{T}_D$ be a plane spanning tree. If T does not contain any twiggly edge, then it is compatible to $\mathcal{S}$. Otherwise, let $e = v_i v_j$ (with $i < j$) be a maximal twiggly edge of T. Note that all twiggly edges in T are non-intersecting, since T is plane. Define

$$V_e^{\uparrow} = (v_k : i < k < j, v_k \text{ is above } e).$$

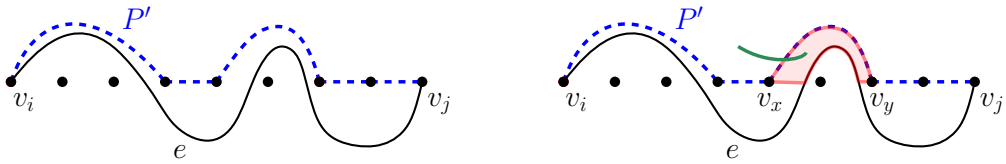

Figure 8.4: *Left:* The (maximal) twiggly edge $e = v_iv_j$ divides the vertices between v_i and v_j into two groups – above and below. The path P' is formed by joining the consecutive vertices lying above e including the vertices of e. *Right:* If an edge of T (green) intersects P', it cannot leave the red region.

Note that $V_e^{\uparrow}$ is not empty. Consider the path $P' = v_i, V_e^{\uparrow}, v_j$ that starts at v_i, ends at v_j and inbetween uses only vertices of $V_e^{\uparrow}$ (see Figure 8.4 (left)). Note that the vertices in P' are in increasing x-order. No edge in P' can intersect the twiggly edge e, since it would have to intersect it twice. Hence, P' lies (strictly) above e. Further note that the interior of the face bounded by $(P' \cup e)$ does not contain a vertex of T.

Next we show that P' cannot intersect T. To this end, let $f \in T$ and assume, for the sake of contradiction, that it intersects an edge v_xv_y of P'. By the properties of simple drawings, f can neither be incident to v_x nor to v_y. However, by construction of P', in order to reach its incident vertices, f has to either (i) intersect v_xv_y twice, or (ii) intersect e, or (iii) intersect a spine edge; see Figure 8.4 (right). (i) yields a contradiction to the properties of simple drawings, (ii) yields a contradiction to the planarity of T, and (iii) yields a contradiction to the maximality of the twiggly edge e.

Hence, we can add P' to T without introducing any crossings. Note that some edges of P' may already be present in T, however, since this insertion creates a cycle, at least one is not. In order to reach a compatible tree, we remove e from this cycle (and potentially more edges until reaching a plane spanning tree again). We created a compatible tree with at least one twiggly edge less and repeating this process, we will eventually reach the spine path $\mathcal{S}$. As we have at most $n-1$ twiggly edges, $\mathcal{F}(\mathcal{T}_D)$ has diameter $O(n)$. □

8.1.3 Strongly c-monotone drawings

A curve is called *c-monotone* (w.r.t. a point x) if every ray emanating from x intersects the curve at most once. A simple drawing is *c-monotone*, if all vertices are drawn along a circle and every edge is a c-monotone curve w.r.t. the center of the circle. A c-monotone drawing is *strongly* c-monotone if for every pair of edges e, e' there is a ray (rooted at the circle center) that intersects neither e nor e'. See Figure 8.5 for an illustration.

In a (strongly) c-monotone drawing, we label the vertices $v_1, v_2, \ldots, v_n$ in cyclic order and denote the center of the circle by c. In the following, we often consider edges and their intersections with rays emanating at c; unless stated

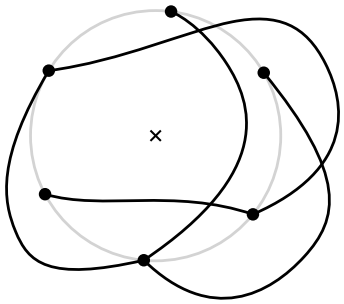

Figure 8.5: Illustration of a strongly c-monotone drawing.

otherwise, any ray is rooted at c and edges are intersected in their relative interiors.

An edge e connecting two consecutive vertices v_i, v_{i+1} is called *cycle edge*. A cycle edge e is called *spine edge* if no ray formed by the center and any other vertex crosses e. All spine edges form the *spine* and any path consisting entirely of spine edges is called *spine path*.

Lemma 8.5. *Let D be a strongly c-monotone drawing of K_n. Then, either all cycle edges in D are spine edges, or D is weakly isomorphic to a monotone drawing.*

Proof. Assume there is an edge $e = v_iv_{i+1}$ of consecutive vertices which does not form a spine edge, i.e., e intersects every ray that is not in the wedge v_icv_{i+1}. The c-monotonicity of edges implies that every edge intersecting the wedge v_icv_{i+1} must intersect every ray in this wedge. This, however, would be in contradiction to the property of *strong* c-monotonicity and hence, the wedge v_icv_{i+1} is not intersected by any edge. This implies the drawing to be weakly[1] isomorphic to a monotone drawing: no ray r in the wedge v_icv_{i+1} is crossed by an edge and hence, we can cut the spine at the intersection with r and stretch the drawing to a monotone drawing. □

Similar to Section 8.1.2, we define *twiggly* edges to be those that intersect a spine edge. A crucial difference to the monotone setting is that an analogue to the relation '$\succ$' (adjusted with respect to the intersection with rays emanating from c) may now be cyclic and hence, we cannot guarantee the existence of a *maximal* twiggly edge anymore. We therefore need a different approach.

For a twiggly edge $e = uw$, let $x_1, \dots, x_k$ be its crossings with the spine (note that these are not vertices of K_n) and assume the labeling to be in such a way that $u, x_1, \dots, x_k, w$ appear in clockwise order. For $i \in \{1, \dots, k\}$ denote the vertex (of K_n) in clockwise order before x_i by x_i^- and the one after x_i by x_i^+. Furthermore, set $u = x_0^-$ and $w = x_{k+1}^+$. Then, for $i \in \{0, \dots, k\}$, we call the edges $x_i^- x_{i+1}^+$ *bumpy* edges (see Figure 8.6 (left)). Note that bumpy edges do

[1] In fact, it even implies strong isomorphism, which, however, is not relevant for the following arguments.

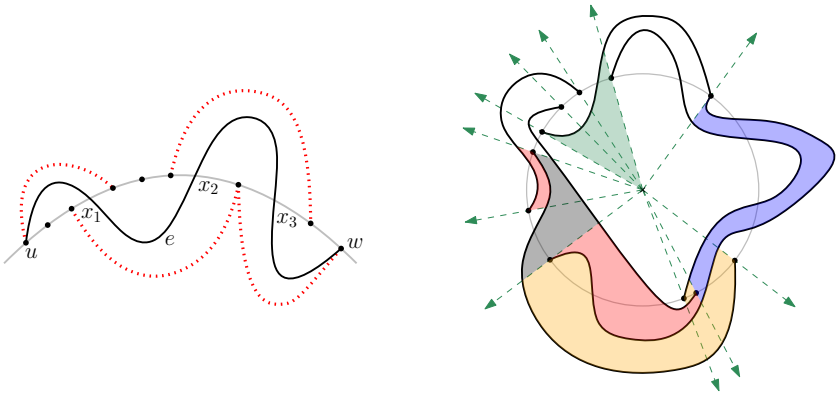

Figure 8.6: *Left:* The red edges are bumpy edges of the twiggly edge e. *Right:* A set of twiggly edges and some corridors; the green is an inner corridor.

not intersect the spine and for every twiggly edge uw there are at least two bumpy edges, namely one incident to u and one incident to w.

We identify a ray r with the angle θ that r forms with the vertical ray (upwards). Two edges e, f are called *neighbours on an interval* $[\theta_1, \theta_2]$, if for any ray $r \in [\theta_1, \theta_2]$ the intersections of e and f with r appear consecutively on r. For a set E of twiggly edges, a *corridor* is a maximal connected region bounded by two neighbouring edges of E (along a maximal interval of angles). Again, we identify corridors by an interval $[\theta_1, \theta_2]$ and usually we speak of corridors defined by the twiggly edges of a plane spanning tree. The *twiggly depth* with respect to a plane spanning tree T of a ray r is the number of twiggly edges of T that r intersects.

We extend our definition of neighbours (along an interval) also to the very inside and very outside by inserting a dummy edge at the circle center and one at infinity. More precisely, an edge e is the neighbor of the circle center c along an interval $[\theta_1, \theta_2]$ if for any ray $r \in [\theta_1, \theta_2]$ the intersection of r and e is closest to c (and furthest in the case of being a neighbor of infinity). We call the corresponding corridors *inner/outer* corridors. See Figure 8.6 (right) for an illustration.

We further remark that for any plane spanning tree T, any corridor $C = [\theta_1, \theta_2]$ of twiggly edges of T begins and ends at a vertex, i.e., the rays at θ_1 and θ_2 hit a vertex.

Lemma 8.6. *For any plane spanning tree T of a strongly c-monotone drawing D of K_n and any corridor C of T with start and end vertex s and t, there is a path P in D from s to t staying entirely in C and not intersecting T. Furthermore, if C is an inner or outer corridor, P does not use any twiggly edge.*

Proof. Let $v_0, v_1, \ldots, v_k, v_{k+1}$ be the vertices of D that are contained in C in cyclic order, where $v_0 = s$ and $v_{k+1} = t$. Furthermore, let $P = v_0, \ldots, v_{k+1}$ be

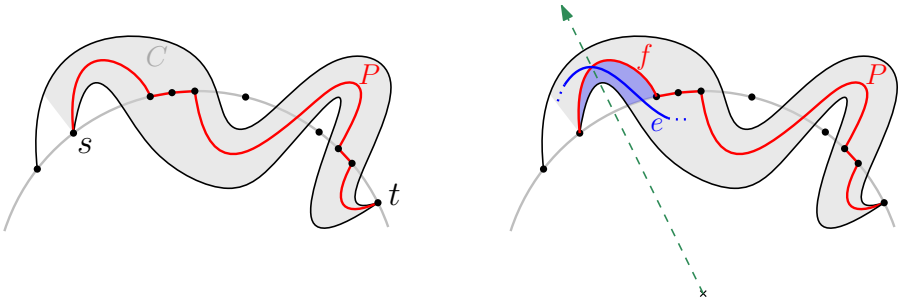

Figure 8.7: Illustration of Lemma 8.6. *Left:* A path constructed inside a corridor. *Right:* If P intersects an edge e of T, the edges bounding the corridor are not neighbours.

the corresponding path that connects consecutive vertices; see Figure 8.7 (left). We claim that P fulfills the desired properties.

First, P stays entirely in C, since any edge leaving C would either violate the properties of simple drawings or strong c-monotonicity.

Second, we argue that P does not intersect T. Assume that there is an edge e of T crossing an edge f of P. Since P connects consecutive vertices in C, we conclude that e intersects the spine, i.e., is a twiggly edge. Hence, at the angle corresponding to the intersection of e and f, the two bounding edges of C are not neighboring twiggly edges in T (see Figure 8.7 (right)); a contradiction.

Lastly, we argue that if P uses a twiggly edge e', which intersects a spine edge uv, then both twiggly edges bounding C must also intersect uv. Assume that one of the bounding edges of C does not intersect uv. Then u or v lie in C. Since P connects consecutive vertices in C, the edge e' would actually use u or v; a contradiction. Hence, if C is an inner or outer corridor, P does not use a twiggly edge. □

Lemma 8.7. *For any strongly c-monotone drawing of K_n, the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter $O(n)$.*

Proof. Let D be a strongly c-monotone drawing of K_n and let T be a plane spanning tree. We show that T can be compatibly transformed to a spine path (by iteratively decreasing its twiggly depth). By Lemma 8.5, we may assume that all n spine edges are present in D. Again, if there is no twiggly edge in T, there is nothing to show.

So let E_{twig} be the set of twiggly edges of T and construct the set $\mathcal{C}$ of all corridors for E_{twig}. Next, for any corridor $C \in \mathcal{C}$ with start and end vertex s and t, we add the path P_C as guaranteed by Lemma 8.6 to T, which preserves planarity by Lemma 8.6.

Then we do not disconnect T when removing E_{twig} now. Indeed, let $e = uw \in E_{\text{twig}}$, then the collection of corridor paths below (and also above) e

connects u and w. So we remove E_{twig} and potentially some further edges until T forms a spanning tree again.

Next we argue that the twiggly depths of all rays must have strictly decreased. Let r be a ray that intersects x previous twiggly edges (i.e., edges from E_{twig}). Then r intersects $x+1$ corridors, two of which are either an inner or outer corridor. By Lemma 8.6 and the properties of c-monotone curves, r intersects at most $x-1$ (new) twiggly edges. Hence, the twiggly depth of any ray decreased by at least one and we recursively continue this process until all rays have twiggly depth 0, in which case T is compatible to a spine path. As we have twiggly depth at most $n-1$ in the beginning, $\mathcal{F}(\mathcal{T}_D)$ has diameter $O(n)$. □

Theorem 2.13 now follows from Lemma 8.3, Lemma 8.4, and Lemma 8.7.

8.2 Special Plane Spanning Trees

In this section, we no longer restrict our drawing anymore, i.e., D will be a *simple drawing* of K_n. Instead we focus on special classes of spanning trees and show that the subgraph $\mathcal{F}(\mathcal{T}_D^*)$ of $\mathcal{F}(\mathcal{T}_D)$ induced by the set of vertices corresponding to stars (0-stars), double stars (1-stars), and twin stars (2-stars) is connected.

The following relation, introduced in [APS+19], will be very useful: Given a simple drawing of K_n with vertex set V and two vertices $g \neq r \in V$, for any two vertices $v_i, v_j \in V \setminus \{g, r\}$, we define $v_i \to_{gr} v_j$ if and only if the edges $v_i r$ and $v_j g$ cross. Aichholzer et al. [APS+19, Lemma 2.1] showed that this relation is asymmetric and acyclic.

We start by showing that stars can always be transformed into each other via a sequence of crossing free edge flips.

Lemma 8.8. *Any two stars in D have distance $O(n)$ in $\mathcal{F}(\mathcal{T}_D^*)$.*

Proof. Given a star T in g (i.e., g is incident to all other vertices of T), we can transform it into a star H in r via a sequence of crossing free edge flips, such that in every step, the graph is a double star with fixed path r,g, in the following way. We label the vertices in $V \setminus \{g, r\}$ such that $v_i \to_{gr} v_j$ implies $i < j$, which is possible due to the acyclicity of $\to_{gr}$ (see Figure 8.8). We iteratively replace an edge gv_i by rv_i starting from $i = n-2$ and continuing in decreasing order. Then all intermediate trees are double stars (with fixed path r,g) and hence, it remains to argue that the flips are compatible, i.e., for $i = n-2, \dots, 1$ the edge gv_i does not cross any edge of the current T. By construction, in any step i, T contains edges of the form (a) rv_j for $j > i$ and (b) gv_k for $k < i$. The edge gv_i cannot cross edges in (a) by the definition of the relation $\to_{gr}$ and also not those in (b) due to the properties of simple drawings. As we need at most $n-2$ steps for the transformation, any two stars have distance $O(n)$ in $\mathcal{F}(\mathcal{T}_D^*)$. □

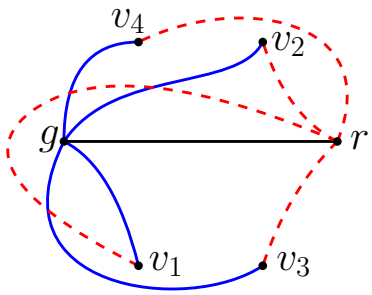

Figure 8.8: Proof of Lemma 8.8. The vertices are labelled in a way conforming to the relation $\rightarrow_{gr}$.

Theorem 2.14 then follows from Lemma 8.8 in combination with the following two lemmata.

Lemma 8.9. *Any double star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.*

Proof. Let T be a double star with fixed path g,r and T' be a star in r'. T can be transformed to a star in r (or g) completely analogous as in the proof of Lemma 8.8. The only difference is that there are additional edges attached to r, i.e., edges of type (a) in Lemma 8.8, which do not interfere. Next, using Lemma 8.8 again, we transform this star in r to T'. This also implies the distance between a double star and a star in $\mathcal{F}(\mathcal{T}_D^*)$ to be $O(n)$. □

Lemma 8.10. *Any twin star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.*

Proof. Let T be a twin star with fixed path g, s, r and T' be a star in r'. Note that all edges in T are incident to g or r and hence, we can add the edge gr and remove gs or rs in order to obtain a double star. Using Lemma 8.9 this double star can be transformed to T'. Thus, a twin star can be tranformed to a star in $O(n)$ steps. □

Chapter 9

Flip Graphs on Arrangements of Pseudocircles

In this chapter we study flip-connectivity of arrangements of pseudocircles with the triangle flip as local transformation. Results and questions about pseudocircles are often motivated by results concerning proper circles or pseudolines. So, let us consider these objects first. We first present the results from a high-level perspective and postpone the precise definitions to the following section.

Ringel [Rin56] studied arrangements of pseudolines with the triangle flip as local transformation. His *homotopy theorem* [Rin56] implies that the corresponding flip graph is connected.

The study of arrangements of pseudocircles was initiated by Grünbaum [Grü72]. A class where flip-connectivity is easy to show is the class of great-pseudocircle arrangements. As they are in bijection to pseudoline arrangements, Ringel's theorem carries over. Note, however, that for great-pseudocircle arrangements, a flip must consist of two simultaneous triangle flips of an antipodal pair of triangles.

Flipping triangles is not enough if disjoint pseudocircles are allowed. In such cases we extend the set of flips by *digon-create* to allow two initially disjoint pseudocircles to start intersecting in a digon and the reverse operation, called *digon-collapse*.

With all the three flips, the flip connectivity of arrangements of (proper) circles is evident since one can shrink all the circles until they have pairwise disjoint interiors. Essentially the same idea works for arrangements of pseudocircles. In this case, however, the fact that a pseudocircle can be shrunk is based on the sweeping lemma of Snoeyink and Hershberger [SH91].

Felsner and Scheucher considered intersecting arrangements of circles as well as pseudocircles. They showed [FS20a, FS20b] that for every $n \in \mathbb{N}$:

(1) The flip graph of intersecting arrangements of n circles is connected.

(2) The flip graph of digon-free intersecting arrangements of n circles is connected.

They also conjecture that their results persist if circles are replaced by pseudocircles:

Conjecture 9.1 ([FS20a, FS20b]). *For every $n \in \mathbb{N}$:*

(1) The flip graph of intersecting arrangements of n pseudocircles is connected.

(2) The flip graph of digon-free intersecting arrangements of n pseudocircles is connected.

As the main result of this chapter we prove part (1) of Conjecture 9.1.

Theorem 2.15. *The flip graph of arrangements of n pairwise intersecting pseudocircles is connected.*

For our proof of Theorem 2.15, we use *cylindrical arrangements*. These are arrangements of pseudocircles in the plane such that the bounded interiors of all the pseudocircles have a common intersection, which we call the *center*. We first show that every cylindrical intersecting arrangement can be flipped into a canonical arrangement by only using triangle flips and without leaving the class of cylindrical arrangements.

Theorem 2.16. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles is connected.*

Showing that every intersecting arrangement $\mathcal{A}$ can be flipped into some cylindrical arrangement then completes the proof of Theorem 2.15. We further study the diameter of flip graphs and obtain asymptotically tight bounds in both settings:

Proposition 2.17. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles has diameter at least $2\binom{n}{3}$ and at most $4\binom{n}{3}$.*

Proposition 2.18. *The flip graph of arrangements of n pairwise intersecting pseudocircles has diameter $\Theta(n^3)$.*

Outline. In Section 9.1, we introduce the necessary concepts and terminology, as well as preliminary results. Section 9.2 is dedicated to the proofs of flip graph connectivity and diameter.

9.1 Preliminaries

A *pseudocircle* is a simple closed curve C which partitions the plane into a bounded region, the interior $\mathrm{int}(C)$, and an unbounded region, the exterior $\mathrm{ext}(C)$. An *arrangement of pseudocircles* is a finite collection of pseudocircles such that every two pseudocircles either are disjoint or they intersect in exactly

one point, where the curves touch, or they intersect in two points, where the curves cross properly.

An arrangement partitions the plane into cells. The 0-dimensional cells or *vertices* are the intersection points (including touchings), the 1-dimensional cells or *edges* are maximal contiguous vertex-free pieces of pseudocircles, and 2-dimensional cells or *faces* are the maximal connected components of the plane after removing all pseudocircles. One can think of the arrangement as a plane graph consisting of vertices, edges, and faces.

Two arrangements are *isomorphic*, i.e., considered the same, if they can be mapped onto each other by a homeomorphism of the plane. Note that we do not require this homeomorphism to be a continuous process (however, in our setting of pseudocircle arrangements, any two isomorphic arrangements can actually be continuously transformed to each other without ever changing the combinatorics[1]). It is noteworthy that our approach of flipping all arrangements to a *canonical arrangement* has the advantage that these isomorphism issues are not of concern. We adopt the common practice and identify arrangements and their isomorphism class.

Every arrangement of circles is an arrangement of pseudocircles and hence, arrangements of pseudocircles generalize arrangements of circles. The generalization is proper since there are *non-circularizable* arrangements, i.e., arrangements of pseudocircles which have no arrangement of circles in their isomorphism class. A thorough study of circularizability was done by Felsner and Scheucher [FS20a].

Throughout this chapter we only consider *simple* arrangements, i.e., there are no touchings and no three curves intersect in a common point. With this convention the cell decomposition of an arrangement of pairwise intersecting pseudocircles is an embedded 4-regular plane graph, possibly with multi-edges. Also recall that all curves are simple, i.e., do not have self-intersections.

A 2-dimensional cell of an arrangement with k edges along its boundary is a *k-face*, a 2-face is a *digon* (some authors call it *empty lens*), and a 3-face is a *triangle*.

It is an instructive exercise to verify that there are exactly four arrangements of three pairwise intersecting pseudocircles (they are shown in Figure 9.1).

Some authors prefer to think of the sphere as the adequate ambient space for arrangements of pseudocircles. Note that in this setting there are only two arrangements of three pairwise intersecting pseudocircles: The arrangements (b), (c), and (d) of Figure 9.1 are isomorphic on the sphere, since they only differ in the choice of the unbounded face. Following [FS20a], we call the arrangement, depicted in Figure 9.1(a), with 8 triangles the *Krupp* arrange-

[1]We remark that in other settings, e.g., arrangements of proper lines, this is not the case. This is referred to as disconnected *realization space*. For instance, Mnëv [Mne88] showed that the realization space for order types is disconnected.

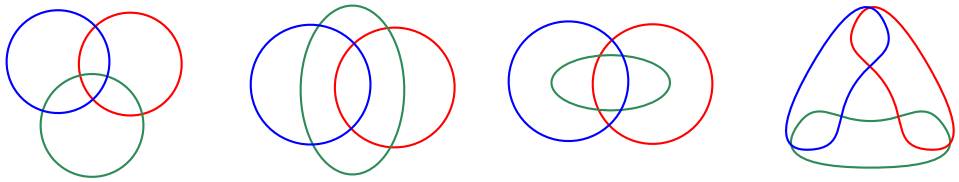

(a) Krupp. **(b)** NonKrupp(4). **(c)** NonKrupp(2). **(d)** NonKrupp(3).

Figure 9.1: The four non-isomorphic arrangements of 3 pairwise intersecting pseudocircles in the plane.

ment and the other ones *NonKrupp*. To make them distinguishable we write NonKrupp(k) to denote the NonKrupp arrangement whose unbounded face has complexity k, e.g. NonKrupp(2) is the arrangement shown in Figure 9.1(c).

Recall that an arrangement of pseudocircles is cylindrical if there is a point p which is contained in the interiors of all the pseudocircles. Among the arrangements of Figure 9.1, the arrangement (d), i.e., the NonKrupp(3), is the only non-cylindrical one.

A subclass of cylindrical arrangements are the great-pseudocircle arrangements. A *great-pseudocircle arrangement* is an arrangement of pseudocircles with the property that every triple of pseudocircles from the arrangement induces a Krupp subarrangement. This class naturally generalizes arrangements of great-circles on the sphere.

9.1.1 Sweeps and Flips

Every two arrangements of (proper) circles can be continuously transformed into each other by moving, expanding, and shrinking circles. The combinatorics of the arrangement changes whenever one circle moves over the crossing of two others, gains two intersections with a circle, or loses its two intersections with a circle.

Snoeyink and Hershberger [SH91] studied continuous transformations of curves and, in particular, of pseudocircles. More precisely, they define the *sweep* of a pseudocircle as a continuous process to expand or shrink the pseudocircle. However, crucially, they also argue that this continuous process can be viewed as a discrete process as the combinatorics of the arrangement changes with one of the following operations:

- a pseudocircle moves over the crossing of two others (*triangle flip*),
- a pseudocircle gains two intersections with a pseudocircle (*digon-create*), or
- a pseudocircle loses its two intersections with a pseudocircle (*digon-collapse*).

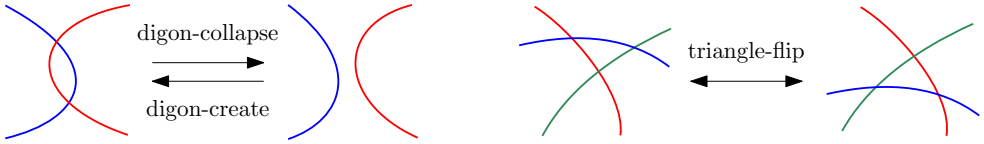

Figure 9.2: An illustration of the three flip operations.

Figure 9.2 depicts the three flip operations. Whenever we speak of *flips* without further specification we refer to these three flips, the term *digon flip* refers to the two flips involving a digon, and otherwise we use the precise term when referring to a specific type of flip.

Snoeyink and Hershberger prove a sweeping lemma (cf. [SH91, Lemma 3.2]) for families of simple curves that pairwise intersect at most twice and are either bi-infinite or closed (i.e., pseudocircles). The flip operations for bi-infinite curves are defined analogously.

Lemma 9.2 (Sweeping Lemma [SH91])**.** *Let $\mathcal{A}$ be an arrangement of pseudocircles or an arrangement of bi-infinite curves that pairwise intersect at most twice. Then $\mathcal{A}$ can be swept starting from any curve C in $\mathcal{A}$ by using three operations: triangle flips, digon-create, and digon-collapse.*

Note that in the context of pseudocircles we sweep towards the *inside* or *outside*, whereas in the context of bi-infinite curves we sweep *upwards* or *downwards*. In other words, Lemma 9.2 guarantees that every pseudocircle C can be continuously expanded or shrunk within $\mathcal{A}$ using the mentioned flips. This can be done until all other pseudocircles are fully contained in the exterior or in the interior of C, respectively. As a consequence, every pseudocircle of an arrangement can be shrunk to an arbitrarily small circle with a finite sequence of flips. Hence, every arrangement can be flipped into the canonical arrangement where all pseudocircles are disjoint. This implies the connectivity of the flip graph of arrangements of pseudocircles. Moreover, the number of flip operations required to flip between two given arrangements is at most $O(n^3)$, which can be seen as follows: Initially a pseudocircle contains at most $n(n-1)$ crossings in its interior and with each triangle flip the number of crossings in the interior decreases. Moreover, each pseudocircle is involved in at most $2(n-1)$ digon flips.

A *lens* in an arrangement is a maximally bounded region in a subarrangement formed by two intersecting pseudocircles. An *arc* is a contiguous subset of a pseudocircle, starting and ending at a vertex of the arrangement.

It will be convenient to have a separate sweeping lemma for lenses. Let $\mathcal{A}$ be an intersecting arrangement of pseudocircles and let Q be (the closure of) a lens bounded by two pseudocircles C_L and C_R. We denote by $L = C_L \cap Q$ and $R = C_R \cap Q$ the two boundary arcs on Q belonging to C_L and C_R, respectively. An *arc of Q* is a maximal connected piece of the intersection of a pseudocircle C with Q, where $C \notin \{C_L, C_R\}$. In other words, an arc of Q is

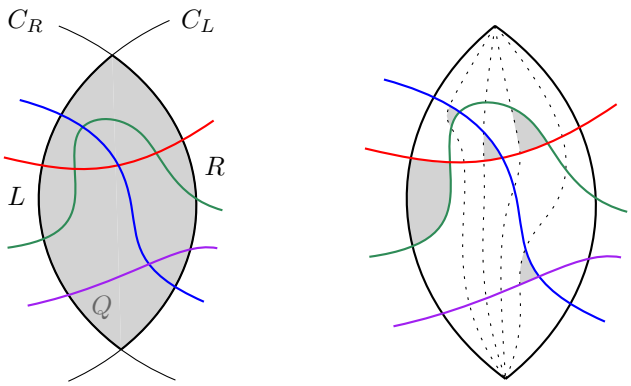

Figure 9.3: Illustration of Lemma 9.3: a lens Q with four transversal arcs and a sweep of Q.

always a contiguous subset of C which has both endpoints on the boundary of Q and whose relative interior lies completely in the interior of Q. If an arc a of Q has both endpoints on L or both endpoints on R, then a forms a lens with L or R respectively. Otherwise the arc has one endpoint on L and one on R, in this case we call the arc *transversal* (see Figure 9.3).

Lemma 9.3. *If Q is a lens and all the arcs of Q are transversal, then, using only triangle flips, L can be swept towards R until the interior of Q does not contain a vertex of the arrangement anymore.*

Proof. We orient all arcs of Q from L to R, and consider the arrangement $\mathcal{A}|_Q$ of these arcs (note that the intersections with L and R do not belong to $\mathcal{A}|_Q$).

Claim (A)**.** *If two oriented arcs a and a' intersect in two points p and q such that p preceedes q on a, then p preceedes q on a'.*

Proof. Suppose for the contrary that q is the first intersection with a on a'. Consider the region T bounded by the subarcs of a and a' from L to q, respectively, and the corresponding part of L; see Figure 9.4(a). Then the arc a' enters the interior of T at p, but cannot leave it anymore to reach R, a contradiction. ■

Claim (B)**.** *The directed graph of the arrangement $\mathcal{A}|_Q$ is acyclic.*

Proof. Suppose that $\mathcal{A}|_Q$ contains a directed cycle. Let Z be a directed cycle such that the enclosed area is minimal. Then Z is the boundary of a face of the arrangement: Indeed, if Z was not a face, there would be a directed cycle contained in Z, contradicting the area minimality of Z; see Figure 9.4(b). Let $v_1, e_1, v_2, \dots e_k, v_{k+1} = v_1$ be the sequence of vertices and edges along Z. From Claim A we know that $k \geq 3$ holds. Let edge e_i be part of arc a_i and let α_i be

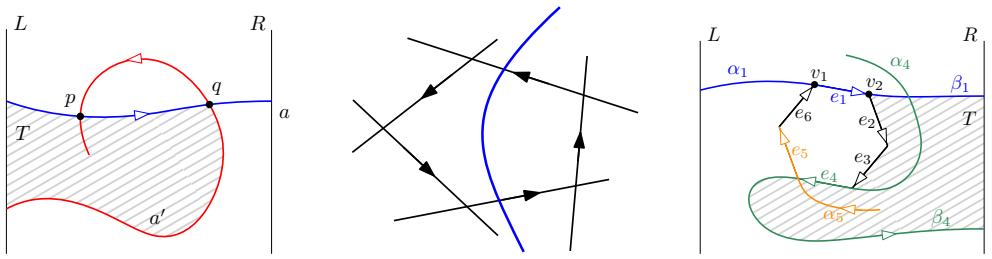

Figure 9.4: Illustration of the proof of Lemma 9.3.

the subarc of a_i starting at L and ending in v_i. Similarly, let β_i be the subarc of a_i starting in v_{i+1} and ending at R.

We will show that for all $j = 2, \ldots, k$ there is an intersection of α_j and β_1. This implies Claim B because the intersection of α_k and β_1 is after v_1 on a_1 and before v_1 on a_k, a contradiction to Claim A.

We argue with induction on j. For $j = 2$, the intersection of α_2 and β_1 is at v_2. For the induction step, given that there is an intersection of α_{j-1} and β_1 we consider the region T in Q whose boundary consists of the directed subpath S from v_2 to v_j on the boundary of Z together with β_1, β_{j-1}, and some piece of R, or in case β_1 and β_{j-1} meet in a point q the boundary consists of S, the arc from v_2 to q on β_1 and the arc from v_j to q on β_{j-1}, see Figure 9.4(c). It is important to note that due to Claim A the interior of Z is not part of T. The final part of α_j is inside T and ends at v_j. Since another intersection of α_j and β_{j-1} is forbidden by Claim A it follows that α_j had to enter region T through β_1. ■

Since the directed graph $\mathcal{A}|_Q$ is acyclic, there is at least one source. We claim that every source gives rise to a triangle with an edge on L that can be flipped. Indeed, let s be a source that is the intersection of two arcs a, a'. Further let x, x' be the intersections of a and a' with L. Since s is a source the subarcs of a and a' from L to s are uncrossed. Moreover, since all arcs in Q are transversal, also the subarc of L from x to x' is uncrossed. Hence, x, x', s form the desired triangle; see also Figure 9.3.

A topological ordering of $\mathcal{A}|_Q$ yields a sequence of triangle flips which can be used to sweep Q with L. This completes the proof of Lemma 9.3. □

For the sake of completeness we kept the standalone proof of Lemma 9.3. However, let us remark that Lemma 9.3 can also be deduced from the sweeping lemma of Snoeyink and Hershberger: The idea is to extend the left boundary curve L of Q and all the arcs of Q to bi-infinite curves. For the extension of each of the curves use two pseudorays – one at each end.

The collection of all pseudorays can be chosen such that they are pairwise disjoint (non-crossing) and do not introduce any further crossings. Further-

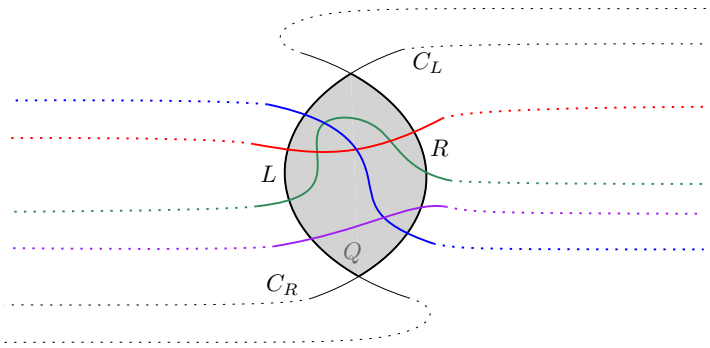

Figure 9.5: All arcs are extended to $-\infty$ and $+\infty$ such that C_L has the property that all transversal arcs are above C_L towards $-\infty$ and below C_L towards $+\infty$.

more, the pseudorays can be chosen such that the sweep with sweep-line L in the resulting arrangement of bi-infinite curves does not require any digon-flips; see Figure 9.5.

Furthermore, it is also possible to add the extending pseudorays in such a way that each pair of the resulting bi-infinite curves has an odd number of crossings and infer Lemma 9.3 from the odd-crossing sweeping lemma of Bokowski et al. [BKPZ18, Lemma 5.2].

9.1.2 Cylindrical Arrangements

An *arrangement of pseudoparabolas* is a finite collection of x-monotone curves defined over a common interval such that every two curves are either disjoint or intersect in two points where the curves cross. The following proposition gives a reversible mapping from cylindrical arrangements to arrangements of pseudoparabolas. The result has been announced by Bultena et al. [BGR99, Lemma 1.3] and a full proof has been given by Agarwal et al. [ANP$^+$04, Lemma 2.11].

Proposition 9.4 ([BGR99, ANP$^+$04]). *A cylindrical arrangement of pseudocircles $\mathcal{C}$ can be mapped to an arrangement of pseudoparabolas $\mathcal{A}$ in an axis-aligned rectangle B such that $\mathcal{C}$ is isomorphic to the arrangement obtained by identifying the two vertical sides of B and mapping the resulting cylindrical surface homeomorphically to a ring in the plane.*

Agarwal et al. [ANP$^+$04, Lemma 2.11] showed that every cylindrical pseudocircle arrangement can be isomorphically represented by a *star-shaped* arrangement, i.e., there exists a point p such that every ray emanating from p intersects all pseudocircles exactly once. Cutting the arrangement along such a ray and stretching it accordingly immediately yields Proposition 9.4.

9.2 Flip Graphs on Pseudocircle Arrangements

The outline of this section is as follows. We prove flip connectivity for cylindrical arrangements of pairwise intersecting pseudocircles in Section 9.2.1 and show in Section 9.2.2 that the corresponding flip graph has diameter at least $2\binom{n}{3}$ and at most $4\binom{n}{3}$. Sections 9.2.3 and 9.2.4 are dedicated to proving flip connectivity for all intersecting arrangements and showing that the diameter of the flip graph is again cubic, where the latter requires several technical arguments.

9.2.1 Proof of Theorem 2.16: Connectivity Cylindrical

In this section we prove flip-connectivity for cylindrical intersecting arrangements by showing that any given arrangement can be flipped to a *canonical* arrangement. We define a canonical intersecting arrangement $\mathcal{B}_n^-$ of pseudoparabolas as follows:

1. On the left, the pseudoparabolas are labelled $C_1, \ldots, C_n$ from top to bottom and on the right, the top to bottom order is the same.

2. For every $1 \leq i \leq n$, the pseudoparabola C_i intersects the other pseudoparabolas first in increasing and then in decreasing order, that is: $C_1, \ldots, C_n, C_n, \ldots, C_1$.

Analogously, $\mathcal{B}_n^+$ denotes the intersecting arrangement of pseudoparabolas in which C_i intersects the other pseudoparabolas first in decreasing and then in increasing order. Closing the pseudoparabolas of $\mathcal{B}_n^-$ and $\mathcal{B}_n^+$ above yields cylindrical intersecting arrangements of pseudocircles, which we call $\mathcal{A}_n^-$ and $\mathcal{A}_n^+$, respectively; see Figure 9.6 for an illustration. Note that every triple of pseudocircles forms a NonKrupp(2) in $\mathcal{A}_n^-$ and a NonKrupp(4) in $\mathcal{A}_n^+$.

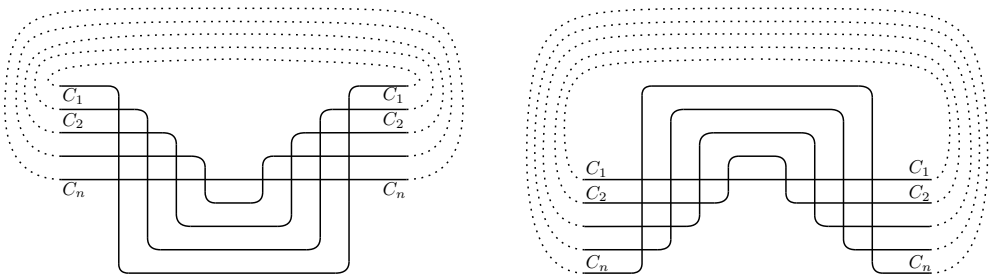

Figure 9.6: The two canonical arrangements $\mathcal{A}_n^-$ (left) and $\mathcal{A}_n^+$ (right).

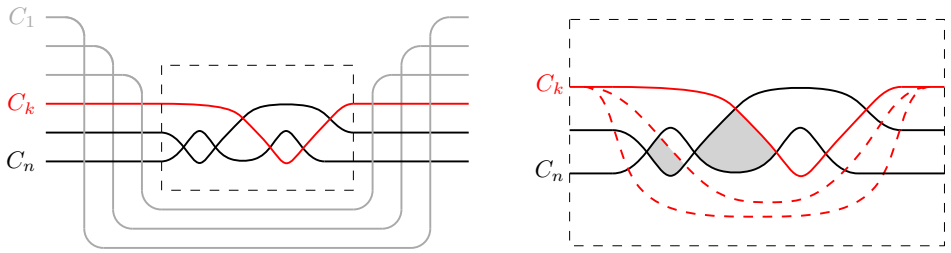

Figure 9.7: An illustration of the proof of Theorem 2.16: flipping the pseudoparabola C_k downwards.

Theorem 2.16. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles is connected.*

Proof. Let $\mathcal{A}$ be an intersecting, cylindrical arrangement of n pseudocircles. Then, using Proposition 9.4, we can represent $\mathcal{A}$ as a pseudoparabola arrangement, in which we label the curves from top to bottom by $C_1, \ldots, C_n$. The idea is to flip the pseudoparabolas downwards one by one in the order of increasing indices. When C_n has been processed the arrangement $\mathcal{B}_n^-$ corresponding to $\mathcal{A}_n^-$ is reached. When flipping C_k downwards the following invariant holds (see Figure 9.7 for an illustration).

1. For every $1 \leq i < k$, the pseudoparabola C_i intersects the other pseudoparabolas first in increasing and then in decreasing order, that is: $C_1, \ldots, C_n, C_n, \ldots, C_1$.

2. For every $i \geq k$, the pseudoparabola C_i first intersects $C_1, ..., C_{k-1}$ in increasing order, then it has the intersections with the pseudoparabolas C_j for $j \in \{k, \ldots, n\} \setminus \{i\}$ in some order, and finally intersects $C_{k-1}, \ldots, C_1$ in decreasing order.

Let $\mathcal{A}_{\geq k}$ denote the subarrangement induced by $C_k, \ldots, C_n$. We claim that in $\mathcal{A}_{\geq k}$ the pseudoparabola C_k can be swept downwards using only triangle flips. Applying Lemma 9.2 to the subarrangement $\mathcal{A}_{\geq k}$, we can sweep C_k downwards using a sequence of triangle, digon-create and digon-collapse flips. However, the latter two flips cannot occur: Since every pair of pseudoparabolas already intersects, a digon-create is impossible. A digon-collapse is impossible because in $\mathcal{A}_{\geq k}$ pseudoparabola C_k is the topmost on the left.

By the invariant, all crossings of $\mathcal{A}_{\geq k}$ lie above all the pseudoparabolas $C_1, \ldots, C_{k-1}$. Therefore, we can perform the flip sequence from $\mathcal{A}_{\geq k}$ also in the original arrangement $\mathcal{A}$. Since we do not flip triangles involving $C_1, \ldots, C_{k-1}$, their intersection orders remain unchanged. Moreover, since in the subarrangement $\mathcal{A}_{\geq k}$ all crossings between $C_{k+1}, \ldots, C_n$ lie above C_k, the

order of intersections in $\mathcal{A}_{\geq k}$ on C_k is $C_{k+1}, \ldots, C_n, C_n, \ldots, C_{k+1}$. From this it follows that the invariant has been preserved.

When C_n has been processed the invariant implies that the canonical arrangement $\mathcal{B}_n^-$ has been reached.

This completes the proof of Theorem 2.16. □

9.2.2 Proof of Proposition 2.17: Diameter Cylindrical

Proposition 2.17. *The flip graph of cylindrical arrangements of n pairwise intersecting pseudocircles has diameter at least $2\binom{n}{3}$ and at most $4\binom{n}{3}$.*

Proof. Concerning the upper bound, following the proof of Theorem 2.16, we can transform any cylindrical arrangement $\mathcal{A}$ of n pairwise intersecting pseudocircles into the canonical arrangement $\mathcal{A}_n^-$. To do so, we iteratively transform the selected pseudocircle C_i by only flipping downwards until all crossings of C_j and C_k with $j, k > i$ are above C_i. For every C_i, there are at most $2\binom{n-i}{2}$ such crossings below C_i and hence, in total, we need at most

$$\sum_{i=1}^{n} 2\binom{n-i}{2} = 2\sum_{i=2}^{n-1}\binom{i}{2} = 2\binom{n}{3}$$

flips to reach $\mathcal{A}_n^-$. Hence, we need at most $4\binom{n}{3}$ flips to transform any two arrangements into each other.

Concerning the lower bound, note that $2\binom{n}{3}$ flips are needed to transform $\mathcal{A}_n^-$ into $\mathcal{A}_n^+$ because every NonKrupp(2) needs to be flipped twice to become a NonKrupp(4). □

9.2.3 Proof of Theorem 2.15: Connectivity Intersecting

For an arrangement $\mathcal{A}$ of n pairwise intersecting pseudocircles, we identify a designated point p and then iteratively expand pseudocircles by triangle flips until all pseudocircles contain p, i.e., until the arrangement is cylindrical. The flip-connectivity then follows from Theorem 2.16.

For convenience we introduce the notation $\mathcal{A} - \{C_1, \ldots, C_k\}$ to denote the arrangement that is obtained from $\mathcal{A}$ by removing the pseudocircles $\{C_1, \ldots, C_k\}$. We say that two pseudocircles C and C' are *parallel* in $\mathcal{A}$ if every vertex of $\mathcal{A} - \{C, C'\}$ lies in $(\text{int}(C) \cap \text{int}(C'))$ or in $(\text{ext}(C) \cap \text{ext}(C'))$.

Theorem 2.15. *The flip graph of arrangements of n pairwise intersecting pseudocircles is connected.*

Proof. Let $\mathcal{A}$ be an intersecting arrangement of n pseudocircles. We show by induction on n that $\mathcal{A}$ can be transformed into a cylindrical arrangement with a finite number of triangle flips. The statement of the theorem then follows from Theorem 2.16.

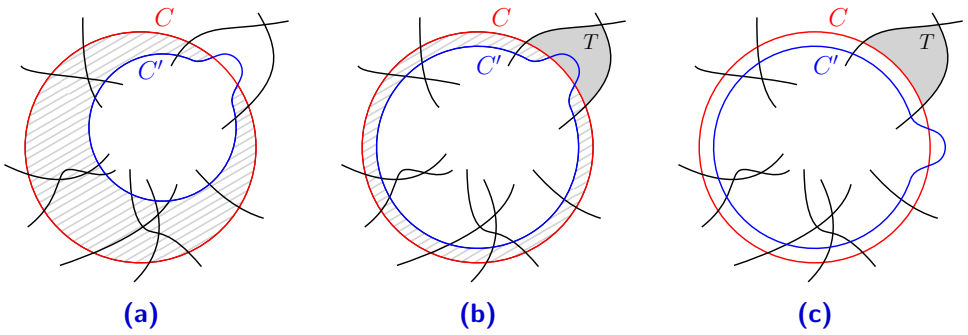

Figure 9.8: (a) C forms a digon with C' that is in $\mathrm{ext}(C) \cap \mathrm{int}(C')$. (b) Expand C' so that it becomes parallel to C. (c) Flip C' so that T becomes also a triangle in $\mathcal{A}$.

The induction base is trivially fulfilled for $n = 2$. For the induction step, we choose p as a point which lies inside a maximum number of pseudocircles. If p lies in the interior of all pseudocircles, then $\mathcal{A}$ is already cylindrical and we are done.

Hence, we may assume that there exists a pseudocircle C which does not contain p in its interior. We show how to expand C until containing p, using only triangle flips. First observe that Lemma 9.2 guarantees the existence of a flip to expand C. Since C already intersects all other pseudocircles, this must be a triangle or digon-collapse flip. As long as there exists a triangle flip expanding C, we perform it and transform $\mathcal{A}$ accordingly.

Suppose now that C does not yet contain p and can only be expanded by collapsing a digon D formed with another pseudocircle C'. Then all remaining pseudocircles must intersect the lens $\mathrm{int}(C) \cap \mathrm{ext}(C')$ transversally (see Figure 9.8(a)). Hence, using Lemma 9.3, we can expand C' until C and C' are parallel (Figure 9.8(b)).

Next consider the arrangement $\mathcal{A}' := \mathcal{A} - \{C'\}$. By the induction hypothesis, $\mathcal{A}'$ can be transformed into a cylindrical arrangement by a finite sequence of triangle flips. We now carefully mimic this flip sequence on $\mathcal{A}$, while maintaining that C and C' are parallel. Suppose that a triangle T in $\mathcal{A}'$ is flipped. If none of the edges of T belongs to C, we can directly apply this triangle flip also in $\mathcal{A}$. If an edge e of T belongs to C and e is crossed by C' then the digon D is located along e. In this case, we apply two triangle flips to C' so that the digon is transferred to one of the two neighboring edges of C as illustrated in Figure 9.8(c), obtaining that e is not crossed by C' (without changing $\mathcal{A}'$). Finally, if e is not crossed by C' then we apply the according triangle flip twice, namely, once for C and once for C'. This concludes the proof. □

9.2.4 Proof of Proposition 2.18: Diameter Intersecting

Extending the notion of parallelism from Section 9.2.3, we call a subset $\mathcal{C}$ of pseudocircles from an intersecting arrangement $\mathcal{A}$ *parallel* if every vertex of the arrangement $\mathcal{A} - \mathcal{C}$ either lies in the exterior of every $C \in \mathcal{C}$ or in the interior of every $C \in \mathcal{C}$. A set of parallel pseudocircles is also called a *bundle*. Note that, for any bundle $\mathcal{C}$ of size k in $\mathcal{A}$, the order of intersections of the $n-k$ pseudocircles from $\mathcal{A} - \mathcal{C}$ is the same for every $C \in \mathcal{C}$. Hence, $\mathcal{A} - \mathcal{C}$ splits $\mathcal{C}$ into $2(n-k)$ *parts*. Moreover, the crossings within the bundle can easily be moved between the parts using only triangle flips; see Figure 9.9.

Let $\mathcal{A}$ be an intersecting arrangement of n pseudocircles. In the proof of Theorem 2.15 we have seen how to transform $\mathcal{A}$ into a cylindrical arrangement with finitely many triangle flips. We now show that $O(n^3)$ flips are sufficient and sometimes necessary. To this end, we first refine the strategy from the proof of Theorem 2.15.

The general situation is that we have a designated point p and a bundle $\mathcal{C}_{[k]} = \{C_1, \ldots, C_k\}$ of pseudocircles, which are simultaneously expanded to eventually contain p. In order to guarantee the cubic bound on the diameter, we need to carefully distribute and move the crossings formed by the bundle in relation to crossings of other pseudocircles with the bundle.

Initial setup. We start with an arbitrary point p from the plane and initialize $\mathcal{D}$ as the set of pseudocircles that already contain p in the interior. We choose C_1 as an arbitrary pseudocircle which is not yet in $\mathcal{D}$, set the bundle size to $k = 1$, set $\mathcal{C}_{[1]} := \{C_1\}$, and proceed with the following steps.

Step 1 (Reset if bundle becomes too large). If $k > n/2$, we reset the process by choosing a new p as an interior point of the bundle $\mathcal{C}_{[k]}$ and set $\mathcal{D} := \mathcal{C}_{[k]}$. Next, we choose C_1 as an arbitrary pseudocircle from $\mathcal{A} - \mathcal{D}$, set the bundle size to $k := 1$, and set $\mathcal{C}_{[1]} := \{C_1\}$. Note that such a reset happens at most once, because, as soon as $|\mathcal{D}| > n/2$ holds, the bundle size k cannot exceed $n/2$ again.

Step 2 (Distribute crossings). If $k \geq 2$ then the k pseudocircles in the bundle $\mathcal{C}_{[k]}$ form $2\binom{k}{2}$ crossings. Also, $\mathcal{C}_{[k]}$ is split by the $n-k$ pseudocircles from $\mathcal{A} - \mathcal{C}_{[k]}$ into $2(n-k)$ parts. The aim is to distribute the crossings of $\mathcal{C}_{[k]}$ such that each part contains at most $\frac{k^2}{n}$ crossings, where $\frac{k^2}{n} \geq \frac{2\binom{k}{2}}{2(n-k)}$ due to $k \leq n/2$. In this case we say that the crossings are *well-distributed*.

We now describe how to maintain the property that crossings are well-distributed when a new pseudocircle C_k enters the bundle $\mathcal{C}_{[k-1]}$. The $2\binom{k-1}{2}$ crossings formed by $\mathcal{C}_{[k-1]}$ are well-distributed among the parts of $\mathcal{C}_{[k-1]}$ that are induced by $\mathcal{A} - \mathcal{C}_{[k-1]}$, i.e., each part contains at most $\frac{(k-1)^2}{n}$ crossings. When C_k enters the bundle, two boundaries between parts of size at most

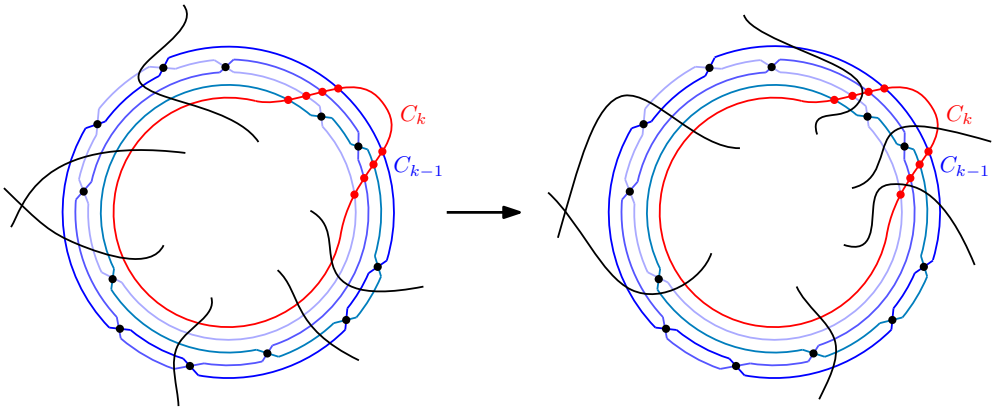

Figure 9.9: An illustration of the distribution of crossings in step 2. *Left:* Well-distributed crossings w.r.t. to the bundle $C_1, \ldots, C_{k-1}$. *Right:* Well-distributed crossings w.r.t. to the bundle $C_1, \ldots, C_k$.

$\frac{(k-1)^2}{n}$ disappear and on the old boundaries, $2k-2$ crossings involving C_k appear, see Figure 9.9. To move crossings from parts of the expanded bundle $\mathcal{C}_{[k]}$ which contain too many crossings into sparser parts of it we use triangle flips.

The cost for moving an excess of one crossing to a part where it complies with the bound is at most $n-k$, i.e., at most $n-k$ triangle flips are needed. Hence, to reach well-distribution for the expanded bundle $\mathcal{C}_{[k]}$, we need at most

$$\left(2\frac{(k-1)^2}{n} + 2k - 2\right) \cdot (n-k) \;<\; \left(2\frac{k^2}{n} + 2k\right) \cdot n \;\leq\; 3kn$$

triangle flips, where we used $k \leq n/2$ for the last inequality.

Step 3 (Expand bundle). Let $\mathcal{A}' = \mathcal{A} - \mathcal{C}_{[k-1]}$. Using Lemma 9.2, we expand C_k in $\mathcal{A}'$ via a sequence of x_k triangle flips until C_k either contains p or forms a digon in $\mathcal{A}'$ with another pseudocircle, which we denote by C_{k+1}. We next show how to mimic this flip sequence in $\mathcal{A}$.

For a triangle T flipped in the sequence in $\mathcal{A}'$, denote by C_k, D, E the three involved pseudocircles. While in $\mathcal{A}'$, the pseudocircle C_k moves over the crossing of D and E, in $\mathcal{A}$ we have to move the entire bundle $\mathcal{C}_{[k]}$ over this crossing. Note that the pseudocircles D and E cross C_k consecutively along C_k in $\mathcal{A}'$. Hence they bound a part P of $\mathcal{C}_{[k]}$, which contains at most $\frac{k^2}{n}$ crossings of $\mathcal{C}_{[k]}$. The strategy is to transfer all crossings of the bundle from the part P to one of its neighboring parts. Each crossing can be transferred by a single triangle flip, so that the total transfer requires at most $\frac{k^2}{n}$ flips.

Next we move the pseudocircles $C_1, \ldots, C_k$ from $\mathcal{C}_{[k]}$ over the crossing of D and E. More precisely, we flip the triangles determined by (C_i, D, E) for all $i \in \{1, \ldots, k\}$. After performing these k flips, we move all previously

transferred crossings back to ensure that each part of $\mathcal{C}_{[k]}$ again contains at most $\frac{k^2}{n}$ crossings, which also requires at most $\frac{k^2}{n}$ flips.

A single flip of C_k in $\mathcal{A}'$ thus maps to a sequence of at most $(k + 2\frac{k^2}{n}) \leq 2k$ triangle flips of $\mathcal{C}_{[k]}$ in $\mathcal{A}$. We call this sequence a *bundle flip*. Since we had x_k triangle flips when expanding C_k in $\mathcal{A}'$, we get a total of at most $2kx_k$ triangle flips to move $\mathcal{C}_{[k]}$ in $\mathcal{A}$.

If the process terminates with C_k containing p in the interior, then also the other pseudocircles $C_1, \ldots, C_{k-1}$ from the bundle contain p in the interior. In this case we add $\mathcal{C}_{[k]}$ to $\mathcal{D}$. Then we pick another pseudocircle C_1 which is not yet in $\mathcal{D}$ arbitrarily, set $k = 1$, and continue with Step 1. Also, if p is in a face which is incident to C_k in $\mathcal{A}'$, then we can move p to the adjacent face inside C_k. This move can also be done in $\mathcal{A}$ where p is moved across the whole bundle. In this case again we add the bundle to $\mathcal{D}$, pick a new C_1 and continue with Step 1. Otherwise we proceed with Step 4.

Step 4 (Add C_{k+1} to bundle). If C_k does not yet contain p and cannot be further expanded by triangle flips, C_k forms a digon with another pseudocircle C_{k+1} from $\mathcal{A} - \mathcal{C}_{[k]}$. Since the digon lies outside of C_k and is incident to C_k, we know that p is not in the digon (otherwise we would have moved p to the other side of C_k). Consider again the arrangement $\mathcal{A} - \mathcal{C}_{[k-1]}$. Using Lemma 9.3, we can expand C_{k+1} in $\mathcal{A} - \mathcal{C}_{[k-1]}$ via triangle flips until it is parallel to C_k. Since none of those flips involves C_k, the same flips can be performed in $\mathcal{A}$. Before those flips, at most $2\binom{n-k-1}{2}$ crossings are located between C_k and C_{k+1}. As each flip reduces this number by one, the expansion process for C_{k+1} requires at most $2\binom{n-k-1}{2} < n^2$ triangle flips. After performing these flips, C_{k+1} is parallel to the entire bundle $\mathcal{C}_{[k]}$. We set $k := k + 1$, $\mathcal{C}_{[k+1]} := \mathcal{C}_{[k]} \cup \{C_{k+1}\}$, and proceed with Step 1.

Analysis. In Step 1, we reset the procedure at most once. Hence, each of Steps 2 to 5 is performed at most twice for each pseudocircle.

Step 2 (well-distribution) is performed for each of the n pseudocircles at most once and each time $O(n^2)$ flips are sufficient. In total, this gives $O(n^3)$ flips.

Step 3 (expand bundle) is a bit more involved. When a bundle has reached p we add the bundle to $\mathcal{D}$ and restart with a new C_1, i.e., a new seed for a bundle. Suppose that bundles of sizes $n_1, n_2, \ldots n_b$ are added to $\mathcal{D}$. Then $n_1 + n_2 + \ldots + n_b \leq n$ holds. During processing a bundle $C_1, \ldots, C_{n_i}$, we perform $(x_1 + \ldots + x_{n_i})$ bundle flips. Whenever we perform a bundle flip, this reduces the number of crossings which lie outside of the pseudocircles $C_1, \ldots, C_{n_i}$. Hence we have $x_1 + \ldots + x_{n_i} \leq 2\binom{n}{2}$ and the number of triangle flips in Step 3 for a bundle of size n_i is bounded by

$$(x_1 + \ldots + x_{n_i}) \cdot 2n_i \leq 2n^2 \cdot n_i.$$

The number of triangle flips in Step 3 for the expansion of all b bundles can therefore be bounded by

$$2n^2 \cdot (n_1 + \ldots + n_b) \leq 2n^3.$$

Step 4 (add to bundle) is again easy. For each of the n pseudocircles, the overall expansion can be performed with at most n^2 flips. For all pesudocircles, this gives $O(n^3)$ flips in Step 4.

Summing up over all steps and multiplying the result by two for the possible case of one reset in Step 1, this completes the proof for the cubic upper bound.

Lower bound. We again consider the two canonical cylindrical arrangements $\mathcal{A}_n^+$ and $\mathcal{A}_n^-$ from Section 9.2.2 and show that their flip-distance in the flip graph of intersecting arrangements is also $2\binom{n}{3}$. Note that the flip graph of cylindrical intersecting arrangements is an induced subgraph of the flip graph of intersecting arrangements. Recall that each triple of pseudocircles in $\mathcal{A}_n^-$ is a NonKrupp(2) while each triple of pseudocircles in $\mathcal{A}_n^+$ is a NonKrupp(4). In the flip graph of arrangements of three pairwise intersecting pseudocircles, the flip distance between NonKrupp(2) and NonKrupp(4) is 2. Hence, at least $2\binom{n}{3}$ triangle flips are needed to get from $\mathcal{A}_n^-$ to $\mathcal{A}_n^+$. This completes the proof of Proposition 2.18.

Chapter 10

Conclusion and further questions

Edge Partitions. In the first part of this thesis, we showed that there exist complete geometric graphs that cannot be partitioned into plane spanning trees. However, there remain several interesting open questions. First of all, it remains wide open how many colors are required to partition any complete geometric graph into plane subgraphs. It would be interesting to settle the leading factor of the n term – a question that was already posed by Bose et al.:

Question 10.1 ([BHRW06, Problem 16])**.** *Does there exist an $\varepsilon > 0$ such that every complete geometric graph K_{2n} can be partitioned into at most $(2-\varepsilon)n$ plane subgraphs?*

Besides the trivial bound that every $K(S)$ with $|S| = 2n$ can always be partitioned into $2n-1$ plane stars, one can also try to approach this question via the result of Bose et al. [BHRW06] that a partition into $2n-k$ plane trees, where k is the size of a largest crossing family on S, is possible. In combination with a result of Pach, Rubin, and Tardos [PRT19], who showed that every point set admits a crossing family of size $n^{1-o(1)}$, Question 10.1 is likely to have a positive answer. In fact, we conjecture that the leading constant is 1, i.e., $\varepsilon = 1$ in Question 10.1.

On the other hand, it would be interesting to investigate *how many* complete geometric graphs cannot be partitioned. In other words, are bumpy wheels the only examples providing a negative answer to Question 3.1, or are there "many" such examples? It also remains open whether $BW_{3,5}$ is the smallest counterexample (with respect to the number of vertices):

Question 10.2. *Does there exist a complete geometric graph with less than 16 vertices that cannot be partitioned into plane spanning trees?*

Recall that Aichholzer et al. [AHK$^+$17] showed that every K_n for $n \leq 10$ can be partitioned into plane spanning trees. Hence, if there is a smaller counterexample, it contains 12 or 14 points.

Question 10.3. *"How many" complete geometric graphs can be partitioned into plane spanning trees?*

This question is wide open. Using the computer program described in Section 3.1, we were unable to find any other counterexamples on point sets sampled uniformly at random from a square. However, given the vast amount of order types (cf. [AAK02]), our program could of course test only a vanishingly tiny fraction of point sets. Further investigating this probabilistic approach, one can consider the following question:

Question 10.4. *What is the probability of a complete geometric graph $K(S)$, where S is sampled uniformly at random from, e.g., a square, to admit a partition into plane spanning trees?*

In the context of partitions into beyond-planar subgraphs, we barely scratched the surface and the obvious next steps would be to investigate non-convex geometric graphs (especially in the context of k-plane subgraphs). One further step in the context of k-quasi-plane subgraphs could be the following question (cf. Theorem 6.5):

Question 10.5. *Let S be a set of n points in general position and $k \geq 3$ be an integer. What is the smallest $c \geq 1$ such that $K(S)$ can be partitioned into $\left\lceil c \cdot \frac{n}{2(k-1)} \right\rceil$ many k-quasi-plane subgraphs?*

We remark that the questions considered in part I of this thesis combine classical graph theoretic questions ("decomposing a graph") with geometric aspects ("into *non-crossing* subgraphs"). Dropping the geometric aspect changes the problem drastically: Tutte [Tut61] and Nash-Williams [NW61] independently obtained necessary and sufficient conditions for an *abstract* graph to admit k edge-disjoint spanning trees. There are many more important questions related to the field of decomposing abstract graphs into certain subgraphs, e.g., Ringel's conjecture [Rin63] and the graceful tree conjecture [Ros67]. The former states that the complete graph K_{2n+1} can be decomposed into copies of any tree with n edges. An affirmative answer for large enough n was given by Montgomery et al. [MPS20].

Concerning the complexity of the corresponding decision problem, Dor and Tarsi [DT97] showed that, given abstract graphs G and H, it is $\mathcal{NP}$ - complete to decide whether G can be partitioned into subgraphs isomorphic to H. Given the fact that not every complete geometric graph can be partitioned into plane spanning trees, it now also makes sense to ask the question concerning the complexity of the decision problem in our geometric setting:

Question 10.6. *What is the complexity of deciding whether a given complete geometric graph K_{2n} can be partitioned into n plane subgraphs / spanning trees?*

Flip graphs. Concerning the second part of this thesis, i.e. flip graphs, there are several questions that remain open.

- *Flipping paths.* In Chapter 7 we proved that edge flips induce a connected flip graph on straight-line paths, if the underlying point set is in generalized double circle position. This covers several important classes of point sets, such as double chains or double circles. For these results we used only Type 1 and Type 2 flips (which can be simulated by Type 1 flips). It is an intriguing question whether Type 3 flips are necessary at all:

  Question 10.7. *Does the connectivity of the flip graph on spanning paths under edge flips depend on whether or not Type 3 flips are allowed?*

  A natural next step towards Conjecture 7.1 would be to adapt our proof of Theorem 2.12 to other classes of point sets, where general spinal point sets are an obvious candidate:

  Question 10.8. *Is the flip graph on spanning paths under edge flips connected for every spinal point set?*

  Most of our results towards Theorem 2.12 hold for spinal point sets in general. The main obstacle is Lemma 7.13.

  A proof for general point sets requires new ideas and seems elusive at the moment.

- *Compatible trees.* In the context of compatible spanning trees, we again proved partial answers towards Question 8.1, namely, we showed connectedness of the compatibility graph for cylindrical, monotone, and strongly c-monotone drawings. Furthermore, we showed connectedness of certain subgraphs, induced by the set of stars, double stars, and twin stars. Again, Question 8.1 remains open in general:

  Question 8.1. *Is the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ connected for every simple drawing D of the complete graph K_n, for all $n \in \mathbb{N}$?*

- *Flipping pseudocircle arrangements.* Concerning the flip connectivity of pseudocircle arrangements we proved part (1) of Conjecture 9.1, showing that triangle flips induce a connected flip graph on intersecting pseudocircle arrangements. As a side-product we also showed connectedness for the subclass of *cylindrical* intersecting arrangements. In both cases we provided an asymptotically tight bound of $\Theta(n^3)$ on the diameter of the flip graph. However, part (2) of Conjecture 9.1 concerning the class of *digon-free* arrangements remains a challenging open question:

Question 10.9. *Is it true that for every $n \in \mathbb{N}$, the flip graph of digon-free intersecting arrangements of n pseudocircles is connected?*

The same question for cylindrical, intersecting arrangements remains open, too.

Lastly, several questions concerning the structure of the flip graphs remain open, e.g., concerning Hamiltonicity, the degree of connectivity, or maximum/minimum degree. Especially in the setting of pseudocircle arrangements, the question concerning the maximum/minimum degree is interesting. Observe that, here, the maximum degree of the flip graph corresponds to the maximum number of triangles among all arrangements and is known to be $\Delta = \frac{4}{3}\binom{n}{2} + O(n)$ [FS21]. The minimum degree δ corresponds to the minimum number of triangles. For intersecting arrangements with digons it is known that $\frac{2n}{3} \leq \delta \leq n-1$ holds and $\delta = n-1$ is conjectured. For intersecting arrangements without digons it is known that $\delta = \lceil \frac{4n}{3} \rceil$ holds for $n \geq 6$ [FRS22, FS21].

Concerning the degree of connectivity, Alves [Alv23] recently showed that the flip graph on *pseudoline* arrangements is $(n-2)$-connected. It is a highly tempting question whether this approach can be adapted also to pseudocircle arrangements:

Question 10.10. *Is the flip graph of intersecting arrangements of n pseudocircles $(n-2)$-connected?*

Bibliography

[ÁAFM+15] 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 *31st European Workshop on Computational Geometry*, pages 57–60, 2015.

[AAH02] Oswin Aichholzer, Franz Aurenhammer, and Ferran Hurtado. Sequences of spanning trees and a fixed tree theorem. *Comput. Geom.*, 21(1-2):3–20, 2002. doi:10.1016/S0925-7721(01)00042-6.

[AAHK06] Oswin Aichholzer, Franz Aurenhammer, Clemens Huemer, and Hannes Krasser. Transforming spanning trees and pseudo-triangulations. *Inf. Process. Lett.*, 97(1):19–22, 2006. doi:10.1016/j.ipl.2005.09.003.

[AAK02] Oswin Aichholzer, Franz Aurenhammer, and Hannes Krasser. Enumerating order types for small point sets with applications. *Order*, 19(3):265–281, 2002. doi:10.1023/A:1021231927255.

[ABD+09] Oswin Aichholzer, Sergey Bereg, Adrian Dumitrescu, Alfredo García Olaverri, Clemens Huemer, Ferran Hurtado, Mikio Kano, Alberto Márquez, David Rappaport, Shakhar Smorodinsky, Diane L. Souvaine, Jorge Urrutia, and David R. Wood. Compatible geometric matchings. *Comput. Geom.*, 42(6-7):617–626, 2009. doi:10.1016/j.comgeo.2008.12.005.

[ABH+18] Oswin Aichholzer, Luis Barba, Thomas Hackl, Alexander Pilz, and Birgit Vogtenhuber. Linear transformation distance for bichromatic matchings. *Comput. Geom.*, 68:77–88, 2018. doi:10.1016/j.comgeo.2017.05.003.

[ABLS15] Greg Aloupis, Luis Barba, Stefan Langerman, and Diane L. Souvaine. Bichromatic compatible matchings. *Comput. Geom.*, 48(8):622–633, 2015. doi:10.1016/j.comgeo.2014.08.009.

[ACK+16] Oswin Aichholzer, Jean Cardinal, Vincent Kusters, Stefan Langerman, and Pavel Valtr. Reconstructing point set order types from radial orderings. *Int. J. Comput. Geom. Appl.*, 26(3-4):167–184, 2016. doi:10.1142/S0218195916600037.

[Ack19] Eyal Ackerman. On topological graphs with at most four crossings per edge. *Comput. Geom.*, 85, 2019. doi:10.1016/j.comgeo.2019.101574.

[ACNS82] M. Ajtai, V. Chvátal, M.M. Newborn, and E. Szemerédi. Crossing-free subgraphs. In Peter L. Hammer, Alexander Rosa, Gert Sabidussi, and Jean Turgeon, editors, *Theory and Practice of Combinatorics*, volume 60 of *North-Holland Mathematics Studies*, pages 9–12. North-Holland, 1982. URL: https://www.sciencedirect.com/science/article/pii/S0304020808734844, doi:10.1016/S0304-0208(08)73484-4.

[AFO+24] Yan Alves Radtke, Stefan Felsner, Johannes Obenaus, Sandro Roch, Manfred Scheucher, and Birgit Vogtenhuber. Flip graph connectivity for arrangements of pseudolines and pseudocircles. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4849–4871, 2024. doi:10.1137/1.9781611977912.172.

[AHK+17] Oswin Aichholzer, Thomas Hackl, Matias Korman, Marc J. van Kreveld, Maarten Löffler, Alexander Pilz, Bettina Speckmann, and Emo Welzl. Packing plane spanning trees and paths in complete geometric graphs. *Inf. Process. Lett.*, 124:35–41, 2017. doi:10.1016/j.ipl.2017.04.006.

[Aic] Oswin Aichholzer. personal communication.

[AIM07] Selim G. Akl, Kamrul Islam, and Henk Meijer. On planar path transformation. *Inf. Process. Lett.*, 104(2):59–64, 2007. doi:10.1016/j.ipl.2007.05.009.

[AKM+22] Oswin Aichholzer, Kristin Knorr, Wolfgang Mulzer, Nicolas El Maalouly, Johannes Obenaus, Rosna Paul, Meghana M. Reddy, Birgit Vogtenhuber, and Alexandra Weinberger. Compatible spanning trees in simple drawings of k_n. In Patrizio Angelini and Reinhard von Hanxleden, editors, *Graph Drawing and Network Visualization - 30th International Symposium, GD 2022, Tokyo, Japan, September 13-16, 2022, Revised Selected Papers*, volume 13764 of *Lecture Notes in Computer Science*, pages 16–24. Springer, 2022. doi:10.1007/978-3-031-22203-0_2.

[AKM+23] Oswin Aichholzer, Kristin Knorr, Wolfgang Mulzer, Johannes Obenaus, Rosna Paul, and Birgit Vogtenhuber. Flipping plane spanning paths. In Chun-Cheng Lin, Bertrand M. T. Lin, and Giuseppe Liotta, editors, *WALCOM: Algorithms and Computation - 17th International Conference and Workshops, WALCOM 2023, Hsinchu, Taiwan, March 22-24, 2023, Proceedings*, volume 13973 of *Lecture Notes in Computer Science*, pages 49–60. Springer, 2023. doi:10.1007/978-3-031-27051-2_5.

[Alv23] Yan Alves Radtke. On the connectivity of the flip-graph of pseudoline arrangements. Bachelor's thesis, 2023.

[ANP+04] Pankaj K. Agarwal, Eran Nevo, János Pach, Rom Pinchasi, Micha Sharir, and Shakhar Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *Journal of the ACM*, 51(2):139–186, 2004. doi:10.1145/972639.972641.

[AOO+22] Oswin Aichholzer, Johannes Obenaus, Joachim Orthaber, Rosna Paul, Patrick Schnider, Raphael Steiner, Tim Taubner, and Birgit Vogtenhuber. Edge partitions of complete geometric graphs. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022, June 7-10, 2022, Berlin, Germany*, volume 224 of *LIPIcs*, pages 6:1–6:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SoCG.2022.6.

[AOV23] Oswin Aichholzer, Joachim Orthaber, and Birgit Vogtenhuber. Towards crossing-free hamiltonian cycles in simple drawings of complete graphs, 2023. arXiv:2303.15610.

[APS+19] 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 2019)*, pages 59:1–59:6, 2019. URL: http://eurocg2019.uu.nl/papers/59.pdf.

[AZ99] Martin Aigner and Günter M. Ziegler. *Proofs from THE BOOK*. Springer, 1999.

[BBD+19] Carla Binucci, Ulrik Brandes, Tim Dwyer, Martin Gronemann, Reinhard von Hanxleden, Marc van Kreveld, Petra Mutzel, Marcus Schaefer, Falk Schreiber, and Bettina Speckmann. *10 Reasons to Get Interested in Graph Drawing*, pages 85–104. Springer International Publishing, Cham, 2019. doi:10.1007/978-3-319-91908-9_6.

[BFRS23] Helena Bergold, Stefan Felsner, Meghana M. Reddy, and Manfred Scheucher. Using sat to study plane hamiltonian substructures in simple drawings. In *39th European Workshop on Computational Geometry*, pages 2:1–2:7, 2023.

[BG20] Ahmad Biniaz and Alfredo García. Packing plane spanning trees into a point set. *Comput. Geom.*, 90:101653, 2020. doi:10.1016/j.comgeo.2020.101653.

[BGR99] Bette Bultena, Branko Grünbaum, and Frank Ruskey. Convex drawings of intersecting families of simple closed curves. In *Proceedings of the 11th Canadian Conference on Computational Geometry (CCCG'99)*, pages 18–21, 1999. URL: http://www.cccg.ca/proceedings/1999/c14.pdf.

[BH09] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009. doi:10.1016/j.comgeo.2008.04.001.

[BHRW06] Prosenjit Bose, Ferran Hurtado, Eduardo Rivera-Campo, and David R. Wood. Partitions of complete geometric graphs into plane trees. *Comput. Geom.*, 34(2):116–125, 2006. doi:10.1016/j.comgeo.2005.08.006.

[BK79] Frank Bernhart and Paul C. Kainen. The book thickness of a graph. *J. Comb. Theory, Ser. B*, 27(3):320–331, 1979. doi:10.1016/0095-8956(79)90021-2.

[BKPZ18] Jürgen Bokowski, Jurij Kovic, Tomaz Pisanski, and Arjana Zitnik. Combinatorial configurations, quasiline arrangements, and systems of curves on surfaces. *Ars Math. Contemp.*, 14(1):97–116, 2018. doi:10.26493/1855-3974.1109.6fe.

[BRUW09] Kevin Buchin, Andreas Razen, Takeaki Uno, and Uli Wagner. Transforming spanning trees: A lower bound. *Comput. Geom.*, 42(8):724–730, 2009. doi:10.1016/j.comgeo.2008.03.005.

[BV11] Prosenjit Bose and Sander Verdonschot. A history of flips in combinatorial triangulations. In Alberto Márquez, Pedro Ramos, and Jorge Urrutia, editors, *Computational Geometry - XIV Spanish Meeting on Computational Geometry, EGC 2011, Dedicated to Ferran Hurtado on the Occasion of His 60th Birthday, Alcalá de Henares, Spain, June 27-30, 2011, Revised Selected Papers*, volume 7579 of *Lecture Notes in Computer Science*, pages 29–44. Springer, 2011. doi:10.1007/978-3-642-34191-5_3.

[CGC] https://cgshop.ibr.cs.tu-bs.de/competition/cg-shop-2022.

[CW09] Jou-Ming Chang and Ro-Yu Wu. On the diameter of geometric path graphs of points in convex position. *Inf. Process. Lett.*, 109(8):409–413, 2009. doi:10.1016/j.ipl.2008.12.017.

[Dav21] James Davies. Box and segment intersection graphs with large girth and chromatic number. *Advances in Combinatorics*, 2021. doi:10.19086/aic.25431.

[Die17] Reinhard Diestel. *Graph Theory*. Springer Publishing Company, Incorporated, 5th edition, 2017.

[DLM19] Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A survey on graph drawing beyond planarity. *ACM Comput. Surv.*, 52(1):4:1–4:37, 2019. doi:10.1145/3301281.

[DLRS10] Jesús De Loera, Jörg Rambau, and Francisco Santos. *Triangulations: Structures for Algorithms and Applications*, volume 25. Springer, 2010. doi:10.1007/978-3-642-12971-1.

[DT97] Dorit Dor and Michael Tarsi. Graph decomposition is np-complete: A complete proof of holyer's conjecture. *SIAM Journal on Computing*, 26(4):1166–1187, 1997. doi:10.1137/S0097539792229507.

[Ede01] Herbert Edelsbrunner. *Geometry and Topology for Mesh Generation*, volume 7 of *Cambridge monographs on applied and computational mathematics*. Cambridge University Press, 2001.

[FRS22] Stefan Felsner, Sandro Roch, and Manfred Scheucher. Arrangements of pseudocircles: On digons and triangles. *CoRR*, abs/2208.12110, 2022. doi:10.48550/arXiv.2208.12110.

[FS20a] Stefan Felsner and Manfred Scheucher. Arrangements of Pseudocircles: On Circularizability. *Discrete & Computational Geometry, Ricky Pollack Memorial Issue*, 64:776–813, 2020. doi:10.1007/s00454-019-00077-y.

[FS20b] Stefan Felsner and Manfred Scheucher. Arrangements of pseudocircles: On circularizability, 2020. arXiv:1712.02149.

[FS21] Stefan Felsner and Manfred Scheucher. Arrangements of Pseudocircles: Triangles and Drawings. *Discrete & Computational Geometry*, 65:261–278, 2021. doi:10.1007/s00454-020-00173-4.

[Gio22] Emeric Gioan. Complete graph drawings up to triangle mutations. *Discret. Comput. Geom.*, 67(4):985–1022, 2022. doi:10.1007/s00454-021-00339-8.

[Grü72] B. Grünbaum. *Arrangements and Spreads*, volume 10 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1972. doi:10.1090/cbms/010.

[GTP21] Alfredo García Olaverri, Javier Tejel Altarriba, and Alexander Pilz. On plane subgraphs of complete topological drawings. *Ars Math. Contemp.*, 20(1):69–87, 2021. doi:10.26493/1855-3974.2226.e93.

[HHN02] M. Carmen Hernando, Ferran Hurtado, and Marc Noy. Graphs of non-crossing perfect matchings. *Graphs Comb.*, 18(3):517–532, 2002. doi:10.1007/s003730200038.

[HHNR05] Michael E. Houle, Ferran Hurtado, Marc Noy, and Eduardo Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs Comb.*, 21(3):325–331, 2005. doi:10.1007/s00373-005-0615-2.

[Hon20] Seok-Hee Hong. Beyond planar graphs: Introduction. In Seok-Hee Hong and Takeshi Tokuyama, editors, *Beyond Planar Graphs, Communications of NII Shonan Meetings*, pages 1–9. Springer, 2020. doi:10.1007/978-981-15-6533-5_1.

[JS08] Björn H. Junker and Falk Schreiber. *Analysis of Biological Networks*. Wiley, 2008.

[KKN04] Seog-Jin Kim, Alexandr V. Kostochka, and Kittikorn Nakprasit. On the chromatic number of intersection graphs of convex sets in the plane. *Electron. J. Comb.*, 11(1), 2004. doi:10.37236/1805.

[Kyn09] Jan Kyncl. Enumeration of simple complete topological graphs. *Eur. J. Comb.*, 30(7):1676–1685, 2009. doi:10.1016/j.ejc.2009.03.005.

[Kyn11] Jan Kyncl. Simple realizability of complete abstract topological graphs in P. *Discret. Comput. Geom.*, 45(3):383–399, 2011. doi:10.1007/s00454-010-9320-x.

[Kyn13] Jan Kyncl. Improved enumeration of simple topological graphs. *Discret. Comput. Geom.*, 50(3):727–770, 2013. doi:10.1007/s00454-013-9535-8.

[Law72] Charles L. Lawson. Transforming triangulations. *Discret. Math.*, 3(4):365–372, 1972. doi:10.1016/0012-365X(72)90093-3.

[Lee89] Carl W. Lee. The associahedron and triangulations of the n-gon. *European Journal of Combinatorics*, 10:551–560, 1989. doi:10.1016/S0195-6698(89)80072-1.

[Lei83] Frank Thomson Leighton. *Complexity Issues in VLSI: Optimal Layouts for the Shuffle-Exchange Graph and Other Networks*. MIT Press, Cambridge, MA, USA, 1983.

[Mne88] N. E. Mnev. *The universality theorems on the classification problem of configuration varieties and convex polytopes varieties*, pages 527–543. Springer Berlin Heidelberg, Berlin, Heidelberg, 1988. doi:10.1007/BFb0082792.

[MPS20] Richard Montgomery, Alexey Pokrovskiy, and Benny Sudakov. A proof of ringel's conjecture, 2020. arXiv:2001.02665.

[Müt22] Torsten Mütze. Combinatorial gray codes-an updated survey. *CoRR*, abs/2202.01280, 2022. arXiv:2202.01280.

[Nis18] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4):52, 2018. doi:10.3390/a11040052.

[NPTZ20] Torrie L. Nichols, Alexander Pilz, Csaba D. Tóth, and Ahad N. Zehmakan. Transition operations over plane trees. *Discret. Math.*, 343(8):111929, 2020. doi:10.1016/j.disc.2020.111929.

[NW61] C. St.J. A. Nash-Williams. Edge-Disjoint Spanning Trees of Finite Graphs. *Journal of the London Mathematical Society*, s1-36(1):445–450, 1961. doi:10.1112/jlms/s1-36.1.445.

[OPG] http://www.openproblemgarden.org/op/partition_of_complete_geometric_graph_into_plane_trees.

[PKK+14] Arkadiusz Pawlik, Jakub Kozik, Tomasz Krawczyk, Michal Lason, Piotr Micek, William T. Trotter, and Bartosz Walczak. Triangle-free intersection graphs of line segments with large chromatic number. *J. Comb. Theory, Ser. B*, 105:6–10, 2014. doi:10.1016/j.jctb.2013.11.001.

[PRT19] János Pach, Natan Rubin, and Gábor Tardos. Planar point sets determine many pairwise crossing segments. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1158–1166. ACM, 2019. doi:10.1145/3313276.3316328.

[PS99] János Pach and József Solymosi. Halving lines and perfect cross-matchings. *Contemporary Mathematics*, 223:245–250, 1999.

[PT97] János Pach and Géza Tóth. Graphs drawn with few crossings per edge. *Comb.*, 17(3):427–439, 1997. doi:10.1007/BF01215922.

[PT06] János Pach and Géza Tóth. How many ways can one draw A graph? *Comb.*, 26(5):559–576, 2006. doi:10.1007/s00493-006-0032-z.

[PW18] Alexander Pilz and Emo Welzl. Order on order types. *Discret. Comput. Geom.*, 59(4):886–922, 2018. doi:10.1007/s00454-017-9912-9.

[Raf88] Nabil H. Rafla. *The Good Drawings D_n of the Complete Graph K_n*. PhD thesis, McGill University, Montreal, 1988. URL: https://escholarship.mcgill.ca/concern/file_sets/cv43nx65m?locale=en.

[Raz08] Andreas Razen. A lower bound for the transformation of compatible perfect matchings. *Proc. of the 24th European Workshop on Comput. Geom.*, 01 2008.

[Rin56] Gerhard Ringel. Teilungen der Ebene durch Geraden oder topologische Geraden. *Mathematische Zeitschrift*, 64:79–102, 1956. doi:10.1007/BF01166556.

[Rin63] Gerhard Ringel. Theory of graphs and its applications. In *Proceedings of the Symposium Smolenice*, 1963.

[Ros67] Alexander Rosa. On certain valuations of the vertices of a graph. *Journal of Graph Theory - JGT*, 1967.

[Sch15] Patrick Schnider. Partitions and packings of complete geometric graphs with plane spanning double stars and paths. Master's thesis, ETH Zürich, 2015.

[Sch16] Patrick Schnider. Packing plane spanning double stars into complete geometric graphs. In *Proc. 32nd European Workshop on Computational Geometry (EuroCG'16)*, pages 91–94, 2016.

[SH91] Jack Snoeyink and John Hershberger. Sweeping arrangements of curves. In Jacob Eli Goodman, Richard Pollack, and William L. Steiger, editors, *Discrete & Computational Geometry: Papers from the DIMACS Special Year*, volume 6 of *DIMACS*, pages 309–349. AMS, 1991. doi:10.1090/dimacs/006/21.

[TCAK19] Hazim Michman Trao, Gek L. Chia, Niran Abbas Ali, and Adem Kilicman. On edge-partitioning of complete geometric graphs into plane trees. 2019. arXiv:1906.05598.

[Tut61] W. T. Tutte. On the problem of decomposing a graph into n connected factors. *Journal of the London Mathematical Society*, s1-36(1):221–230, 1961. doi:10.1112/jlms/s1-36.1.221.

[Wag36] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936. URL: http://eudml.org/doc/146109.

Appendix A

Source code

The source code that we used for partitioning complete geometric graphs (part I of this thesis) is available as github repository and can be downloaded from:

`https://github.com/jogo23/edge_partitions_cgg`

To obtain the results presented in this thesis, we used the code at the following commit id:

`d338a54abf3d0b0bfe8420a41dc134f98eea4963`

Appendix B

Additional Figures

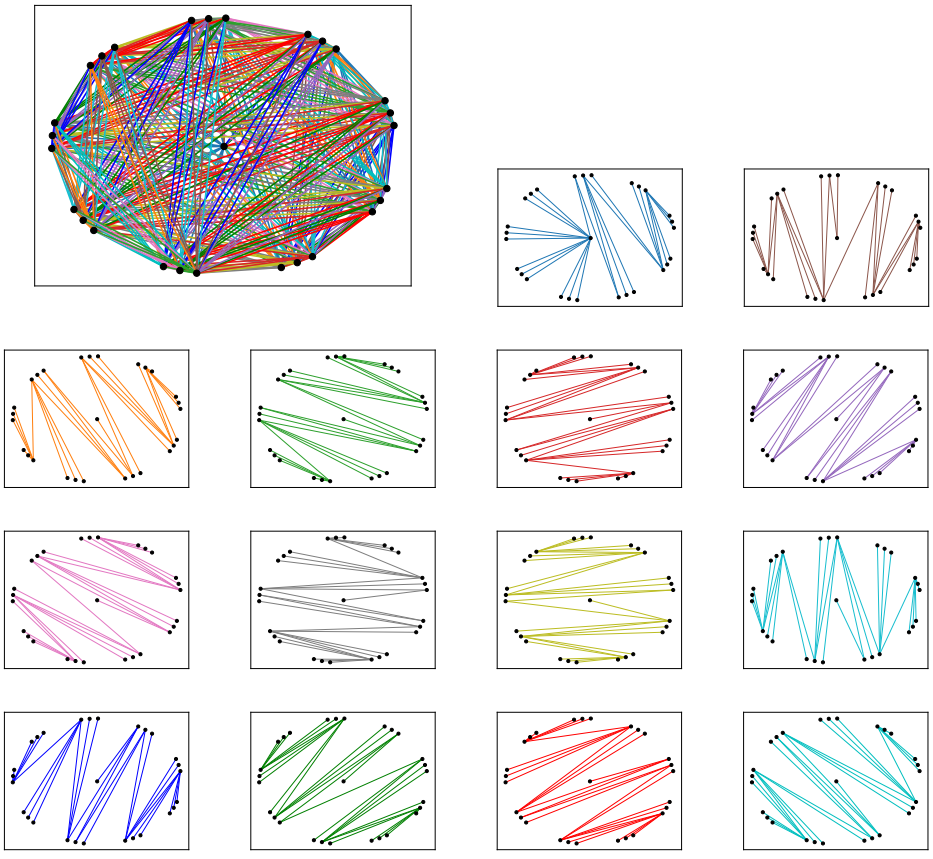

Figure B.1: Full partition of $BW_{9,3}$ into 14 plane spanning trees according to the construction in case 1 of Lemma 3.11 (generated by the computer assisted ILP).

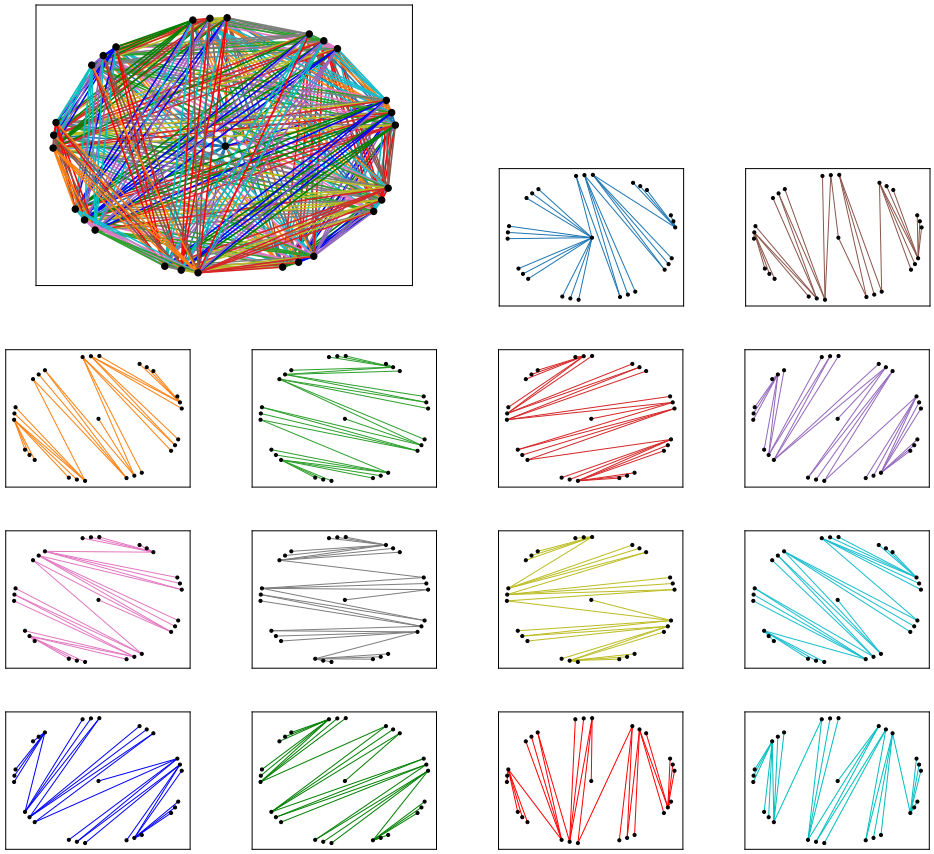

Figure B.2: Full partition of $BW_{9,3}$ into 14 plane spanning trees according to the construction in case 2a of Lemma 3.11 (generated by the computer assisted ILP).

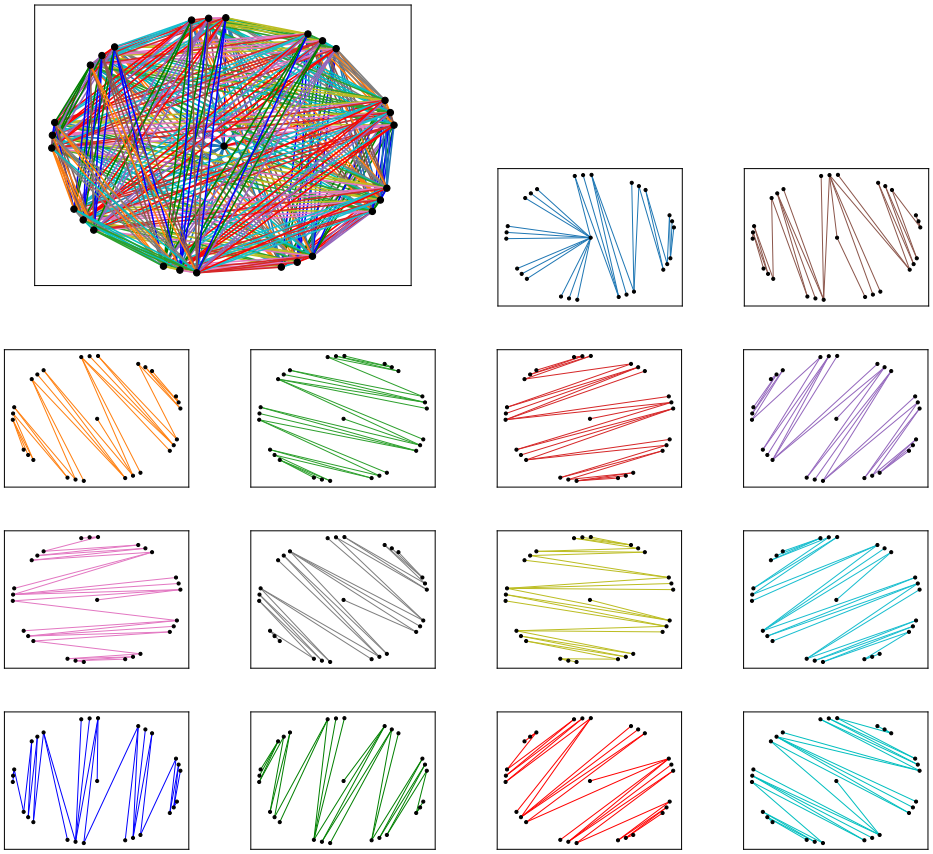

Figure B.3: Full partition of $BW_{9,3}$ into 14 plane spanning trees according to the construction in case 2b of Lemma 3.11 (generated by the computer assisted ILP).

Zusammenfassung

Wir untersuchen zwei fundamentale Themen aus dem Bereich der kombinatorischen Geometrie: planare Graphen und Rekonfigurationsprobleme.

Partitionierung der Kanten von vollständigen geometrischen Graphen. Dieser Teil der Arbeit ist einer langjährigen Frage gewidmet, die zuerst von Bose, Hurtado, Rivera-Campo und Wood [BHRW06] publiziert wurde: Können die Kanten von jedem vollständigen geometrischen Graphen mit einer geraden Anzahl Knoten in planare Spannbäume partitioniert werden? Positive Antworten sind bekannt für Spezialfälle, nämlich wenn die Knotenmenge in konvexer oder regulärer Rad-Lage ist [BHRW06, AHK+17, TCAK19].

Wir zeigen, dass eine solche Partitionierung nicht immer möglich ist und darüberhinaus, dass dies auch für Partitionierungen in planare Teilgraphen nicht immer möglich ist. Alle Konstruktionen basieren auf Punktmengen in Rad Lage, die wir genauer untersuchen und Charakterisierungen angeben unter welchen Bedingungen Partitionierungen existieren oder nicht.

Außerdem untersuchen wir das Problem in allgemeineren Settings jenseits planarer Teilgraphen und zeigen Schranken an die Anzahl der benötigten Teilgraphen für Partitionierungen in k-planare und k-quasiplanare Teilgraphen.

Flip-Graphen. Eine natürliche Art und Weise die Struktur von Rekonfigurationsproblemen zu untersuchen ist über sogenannte *Flip-Graphen*: für eine Menge $\mathcal{X}$ von Objekten und eine lokale Flip-Operation f, hat der Flip-Graph auf $\mathcal{X}$ unter f einen Knoten für jedes Element in $\mathcal{X}$ und zwei Knoten sind benachbart genau dann, wenn die korrespondierenden Objekte mit einem Flip ineinander überführt werden können. Die erste fundamentale Frage, die üblicherweise von Bedeutung ist, ist die Frage, ob ein Flip-Graph zusammenhängend ist. Wir untersuchen die folgenden drei Konstellationen:

- **Flip-Graphen für planare Pfade.** Wir zeigen, dass Kanten-Flips einen zusammenhängenden Flip-Graphen induzieren für die Spezialfälle, dass die Knoten der Pfade in Rad- oder generalisierter Doppelkreis Lage sind.
- **Kompatible Bäume.** Wir zeigen, dass Kompatibilitäts-Flips einen zusammenhängenden Flip-Graphen induzieren für die Spezialfälle, dass die zugrundeliegende Zeichnung zylindrisch, monoton oder stark c-monoton ist.
- **Flip-Graphen für Pseudokreise.** Wir zeigen, dass Dreiecks-Flips einen zusammenhängenden Flip-Graphen auf Arrangements von paarweise schneidenden Pseudokreisen induzieren.