DISS. ETH NO. 28594

Maximum Clique in Generalisations of Disk Graphs and Plane Geometric Graphs on Degenerate Point Sets

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZÜRICH

(Dr. sc. ETH Zürich)

presented by

Nicolas Grelier
Ingénieur diplômé, IMT Atlantique

born on 16.03.1995

accepted on the recommendation of

Prof. Dr. Emo Welzl, examiner
Dr. Marthe Bonamy, co-examiner
Dr. Michael Hoffmann, co-examiner
Prof. Dr. Wolfgang Mulzer, co-examiner

2022

Abstract

This thesis deals with graphs having geometric representations. On one hand we consider graphs whose vertices can be mapped to geometric objets in an Euclidean space (for instance disks in the plane), such that two vertices are adjacent if and only if the corresponding objects intersect. Those are called “intersection graphs”, and if all objects are constrained to be, e.g. disks, then they are referred to as “disk graphs”. On the other hand we study graphs that can be represented in the plane such that vertices are mapped to points and edges to straight-line segments, such that no two edges cross. Those are called “plane geometric graphs”. In this thesis, we also have additional conditions. We consider for instance convex partitions, for which the union of the faces is equal to the convex hull of the points and each bounded face is convex. As a special case, we also study triangulations, for which we additionally require that every bounded face be a triangle.

For intersection graphs, we study the problem of finding a maximum clique in disk-like intersection graphs. In 1990, a seminal paper by Clark, Colbourn and Johnson showed that maximum clique can be solved in polynomial time in unit disk graphs. However, the complexity of maximum clique in disk graphs is still unknown. Recently,

Bonamy et al. showed the existence of an EPTAS for maximum clique in disk graphs. This leads to the following questions: Are there superclasses of unit disk graphs in which maximum clique can be solved in polynomial time? Are there superclasses of disk graphs for which there is an EPTAS? Are there related classes for which we can show NP-hardness? Concerning the first question, we show that maximum clique can be solved in polynomial time in intersection graphs of translates of a fixed bounded convex set. Furthermore, we define a superclass C of both unit disk graphs and interval graphs, where C is defined as the intersection graph class of some specified sets, for which there exists a polynomial time algorithm. For the second question, we prove the existence of an EPTAS for homothets of a fixed bounded and centrally symmetric convex set. We also give partial results toward showing the existence of an EPTAS for intersection graphs of convex pseudo-disks. Finally, for the third question, we show that maximum clique is NP-hard, and even APX-hard, in intersection graphs of unit disks and axis-parallel rectangles.

Concerning triangulations and convex partitions, we study two problems that were previously only considered under the assumption that no three of the n input points are on a line. For triangulations, we extend a result by Wagner and Welzl and show that the bistellar flip graph is $(n-3)$-connected. For convex partitions, we provide the first approximation algorithms for computing convex partitions with as few faces as possibles when three points or more may lie on a line. In particular, we give an $\mathcal{O}(\log(OPT))$-approximation algorithm running in time $\mathcal{O}(n^8)$, where OPT denotes the size of a minimum solution. We additionally provide an $\mathcal{O}(\sqrt{n}\log(n))$-approximation algorithm running in time $\mathcal{O}(n^2)$. We also show that minimising the number of faces is NP-hard.

Résumé

Cette thèse a pour sujet les graphes qui admettent des représentations géométriques. D'une part, nous considérons des graphes dont on peut associer les sommets avec des objets géométriques dans un espace euclidien (par exemple des disques dans le plan), de sorte que deux sommets soient adjacents si et seulement si les objets correspondants s'intersectent. Ces graphes sont appelés "graphes d'intersection", et si tous les objets sont, par exemple, des disques, on parle alors de "graphes de disque". D'autre part, nous étudions des graphes qui peuvent être représentés dans le plan, de sorte que les sommets soient représentés par des points, les arêtes par des segments, et que les arêtes ne se croisent pas. Ceux-ci sont appelés "graphes géométriques planaires". Dans cette thèse, nous posons des restrictions additionnelles. Nous considérons par exemple les partitions convexes, pour lesquelles chaque face bornée doit être convexe. Nous étudions aussi un cas particulier, les triangulations, pour lesquelles il faut également que chaque face bornée soit un triangle.

Concernant les graphes d'intersection, nous considérons le problème de la clique maximale dans des graphes d'intersection d'objets similaires à des disques. En 1990, un papier fondateur de Clark, Colbourn et Johnson montra que le problème de la clique maximum

peut être résolu en temps polynomial dans les graphes de disque-unité. Cependant, la complexité du problème de la clique maximum dans les graphes de disque reste inconnue. Récemment, Bonamy et al. ont montré l'existence d'un EPTAS (schéma d'approximation en temps polynomial) pour la clique maximale dans les graphes de disque. Les questions suivantes viennent naturellement : existe-t-il des superclasses des graphes de disque-unité pour lesquelles le problème de la clique maximale peut être résolu en temps polynomial ? Existe-t-il des superclasses des graphes de disque pour lesquelles le problème peut être approximé par un EPTAS ? Existe-t-il des classes similaires pour lesquelles l'on peut montrer que le problème est NP-difficile ? Pour la première question, nous montrons que le problème de la clique maximale peut être résolu en temps polynomial dans les graphes d'intersection de translatés d'un ensemble convexe fixé. De plus, il existe une superclasse C des graphes de disque et des graphes d'intervalle, où C est définie comme la classe de graphes d'intersection de certains ensembles, pour laquelle il existe un algorithme polynomial. Concernant la seconde question, nous montrons l'existence d'un EPTAS pour les homothètes d'un ensemble fixé, borné et centralement symétrique. Nous donnons aussi des résultats partiels au sujet de l'existence d'un EPTAS pour les graphes d'intersection de pseudo-disques convexes. Enfin, concernant la troisième question, nous montrons que le problème de la clique maximum est NP-difficile, et même APX-difficile, dans les graphes d'intersection de disques unités et de rectangles dont les côtés sont parallèles aux axes.

Au sujet des triangulations et des partitions convexes, nous étudions deux problèmes qui avaient été préalablement considérés seulement sous la condition que pour tout triplet de points, les trois points ne soient pas alignés. Concernant les triangulations, nous étendons un résultat de Wagner et Welzl et montrons que le graphe des flips bistellaire est $(n-3)$-connecté. Pour les partitions convexes, nous donnons les premiers algorithmes d'approximation pour trouver des partitions convexes avec aussi peu de faces que possible lorsque trois points ou plus peuvent être alignés. En particulier, nous donnons un algorithme réalisant une $\mathcal{O}(\log(OPT))$-approximation en temps $\mathcal{O}(n^8)$, où OPT désigne la taille d'une solution minimum. De plus, nous donnons un algorithme réalisant une $\mathcal{O}(\sqrt{n}\log(n))$-approximation en temps $\mathcal{O}(n^2)$. Nous montrons aussi que minimiser le nombre de faces est NP-difficile.

Acknowledgments

My deepest gratitude goes to Michael Hoffmann for his invaluable supervision. He was always there to advise me, and I am extremely fortunate to have been guided by someone so involved in his role. His high expectations pushed me to reach the best results I could, and write this thesis of which I can now be proud.

My sincere thanks to Emo Welzl for accepting me in his group. He also encouraged me to aim for the best. It was a pleasure to work with him on triangulations. I keep very fond memories of when we started working on the problem, all the more so as it was my first in-person collaborative research since the beginning of the pandemic.

I want to thank the other members of my committee. My sincere appreciation goes to Marthe Bonamy, who is also one of the authors of the paper that got me started to the field of maximum clique in disk graphs. I am especially thankful to Wolfgang Mulzer, who is one of my co-authors on two papers that are not part of this thesis. I want to thank him again for welcoming me in his group at FU Berlin during a research travel, despite it being shortened due to the pandemic.

I thank Édouard Bonnet for inviting me twice to Lyon, for research trips that were very fruitful. I also thank him for the two nice papers, including one with Marthe Bonamy, that he co-wrote and which got me interested in the field.

My gratitude goes to Hung Hoang, who shared his office with me for all those years. He was a great help, especially when I started. I also thank Nicolas El Maalouly who joined us later, and who accepted to have his name mispronounced so that we could tell one Nicolas from the other.

I would like to thank every one who guided on my mathematical journey and help me start my doctoral studies. In particular, my deepest gratitude goes to Lionel Chaussade, François Sauvageot, Bastien Pasdeloup, Vincent Gripon, Zhiyi Huang, Rémi de Joannis de Verclos and Ross Kang.

I want to thank all of my co-authors of papers written during and before my doctoral studies whom I have not mentioned yet. Those are Jean-Charles Vialatte, Carlos Eduardo Rosar Kós Lassance, Elsa Dupraz, François Pirot, Saeed Gh. Ilchi, Tillmann Miltzow, Shakhar Smorodinsky, Jonas Cleve, Kristin Knorr, Maarten Löffler, Daniel Perz, Helena Bergold, Daniel Bertschinger and Patrick Schnider.

I would like to thank again all members of our group with whom I had the pleasure to meet and work, and whom I have not yet mentioned: Bernd Gärtner, Luis Barba, Manuel Wettstein, Malte Milatz, Jerri Nummenpalo, Ahad Noori Zehmakan, Tim Taubner, Meghana M. Reddy and Simon Weber. Also, my sincere thanks to Andrea Salow for all the help with the administrative tasks, and of course for organising our weekly lunch meetings.

Finally, I thank my family for their support, with special thanks to my mother Sylvie and my grandmother Yvonne, and my heartfelt thanks to Lucie.

Zürich, July 2022 *Nicolas Grelier*

Contents

CHAPTER 1

Introduction

The topics studied in this thesis are at the intersection between graph theory, geometry and complexity theory. Graphs are commonly used to represent relations between different objects. The graphs we consider here originate from some geometric considerations. Those are twofold. On one hand, we consider intersection graphs of geometric objects. There is one vertex for each object, and two vertices are adjacent if and only if the two objects intersect. An intersection graph therefore summarises the relations between the objects with respect to intersections. On the other hand, we study plane geometric representations of graphs, where each vertex is mapped to a point in an Euclidean space, a (straight-line) segment connects the two points corresponding to a pair of adjacent vertices, and no two segments cross. Roughly speaking, in the first scenario, the geometry related to the vertices is interesting, whereas in the second scenario the edges are what matters. As mentioned, we refer to the first type of graphs as *intersection graphs* and to the second type as *plane geometric graphs*. Another interesting topic that mixes graph theory and geometry is the one of graph drawing. In this setting, each vertex is mapped to a point in the plane, but now edges are represented as simple curves, instead of being constrained to segments. However, this topic is not covered in this thesis.

Concerning intersection graphs, the geometric objects we consider in this thesis include unit disks, disks with arbitrary radii, and axis-parallel rectangles among others. In particular, all objects considered are convex. Convexity also plays a fundamental role in the chapters of the thesis about geometric graphs. Indeed, a part of our results concerns convex partitions: plane geometric graphs where all bounded faces are convex, and the union of the bounded faces is equal to the convex hull of the vertex set (seen as a point set).

We study these topics through the lens of complexity theory. We consider maximisation and minimisation problems, and try to find efficient algorithms to solve them, or try to show that none exists. In our language, "efficient" means that the algorithms run in polynomial time. Under the assumption that P is not equal to NP, we show that some problems cannot be solved, or even approximated, in polynomial time.

1.1 Definitions and notations

1.1.1 Graph notations

Let G be a simple graph. We denote by $V(G)$ the vertex set of G. We say that two vertices are *adjacent* if there is an edge between them, otherwise they are *independent*. If v' is adjacent to v, for v and v' in $V(G)$, we also say that v' is a *neighbour* of v. For a vertex v in $V(G)$, the set $\mathcal{N}(v)$ denotes its *neighbourhood*, i.e. the set of vertices adjacent to v.

Let H be an induced subgraph of G, meaning that $V(H)$ is a subset of $V(G)$ and two vertices in $V(H)$ are adjacent in H if and only if they are adjacent in G. We denote by $G \setminus H$ the subgraph induced by $V(G) \setminus V(H)$. We denote by $\overline{G}$ the *complement* of G, which is the graph with the same vertex set, but where edges and non-edges are interchanged. A *clique* is a set of pairwise adjacent vertices. An *independent set* is a set of pairwise non-adjacent vertices. A graph is *bipartite* if its vertex set can be partitioned into two independent sets. A graph is *cobipartite* if its complement is a bipartite graph. The *clique number*, $\omega(G)$, is the size of a maximum clique. The *independence number*, $\alpha(G)$, is the size of a maximum independent set.

1.1.2 Geometric notations

Throughout the thesis we only consider Euclidean spaces with the Euclidean distance. Let P be a point set in $\mathbb{R}^d$, where d is a positive integer. We say that P is in *general position* if no $k+2$ points lie on an affine subspace of dimension k, where $k < d$. If a point set is not in general position, we say that it is in *degenerate position.* The words "general position" are the usual way of qualifying these types of points sets, but note that they can be misleading: Assuming that a point set is in general position is actually a significant restriction for many problems. Some results of this thesis illustrate perfectly this fact.

The *convex hull* of P, *conv*(P), is the intersection of all convex sets that contain P. A point p in P is said to be *extreme* if P and $P \setminus \{p\}$ do not have the same convex hulls. We say that P is in *convex position* if all points are extreme. A point p in P is an *interior point* or *inner point* if it lies in the interior of *conv*(P). Observe that a point may be both not extreme and not an inner point. If this occurs, then P is in degenerate position. A point set in $\mathbb{R}^2$ in convex position is also in general position, but this does not necessarily hold in higher dimensions.

A *geometric graph* is an embedding of a graph in the plane, where each edge is realised as a straight-line segment. If the edges only meet at vertices, we say that the embedding is *plane*. A *convex partition* is a plane geometric graph, where each bounded face is convex, each vertex has positive degree, and the union of the bounded faces is equal to the convex hull of the vertex set. A *triangulation* is a convex partition where each bounded face is a triangle.

Let p and p' be two points in $\mathbb{R}^d$. We denote by (p, p') the line going through them, and by $[p, p']$ the line segment with endpoints p and p'. We denote by $d(p, p')$ the distance between p and p'. We denote by O_d (or simply by O when it is clear from the context) the origin in $\mathbb{R}^d$. When $d = 2$, we denote by Ox and Oy the x and y-axis, respectively. For $d = 3$, we denote by xOy the xy-plane.

Definition 1.1. Let S be a family of subsets of $\mathbb{R}^d$. We denote the *intersection graph* of S by $G(S)$. It is the graph whose vertex set is S and where there is an edge between two vertices if and only if the corresponding sets in S intersect.

One can easily observe that all graphs are intersection graphs of some connected sets in $\mathbb{R}^3$. Let G be a graph, and let P be a set of $|V(G)|$ points in convex and general position in $\mathbb{R}^3$. Let us consider any bijective map ϕ from $V(G)$ to P. For $v \in V$, we define the set s_v as the union of all segments $[\phi(v), (\phi(v) + \phi(v'))/2]$, where v' is a neighbour of v. Now observe that s_v and $s_{v'}$ intersect if and only if v and v' are adjacent. In this thesis, we consider only some constrained types of geometrical objects, therefore all intersection graph classes that we consider have forbidden graphs.

For a given family of objects, say unit disks, a graph is said to be a *unit disk graph* if it is the intersection graph of some unit disks in the plane. We do likewise with all objects, e.g. disks, axis-parallel rectangles, segments, and so on. Note that any of those are abstract graphs. For a unit disk graph G, a *representation* of G is a set of unit disks U such that G is the intersection graph of U. Likewise, we speak of representations of disk graphs, segment graphs, and so on. A representation is generally not unique.

A family of closed sets in the plane is a *pseudo-disk arrangement* if the boundaries of any two sets intersect at most twice. In this thesis, we consider *convex pseudo-disk graphs*, i.e. intersection graphs of arrangements of pseudo-disks, where each set in the family is convex.

We name a ball of dimension d living in $\mathbb{R}^d$ a *d-ball*. Therefore 1-ball graphs are interval graphs and 2-ball graphs are disk graphs. When d is omitted we implicitly assume $d = 3$, i.e. a ball is a 3-ball.

We sometimes consider intersection graphs of two types of objects, for instance unit disks and axis-parallel rectangles. Therefore when we speak of an *intersection graph of unit disk and axis-parallel rectangles*, we are referring to a graph that admits a representation where the sets in the plane can be unit disks and axis-parallel rectangles. In general, the class of intersection graphs of two objects A and B is a strict superclass of the union of the classes of A-graphs and B-graphs.

1.1.3 Complexity notation

A *polynomial-time approximation scheme* (PTAS) for a maximisation problem is an algorithm which takes, together with its input of size n, a parameter $\varepsilon > 0$ and outputs in time $n^{f(\varepsilon)}$ a solution of value at least $(1-\varepsilon)$OPT, where OPT is the optimum value. An *efficient* PTAS (EPTAS) is the same but has running time $f(\varepsilon)n^{\mathcal{O}(1)}$. Note that the existence of an EPTAS, for a problem in which the objective value is the size of the solution k, implies an FPT algorithm in k, by setting ε to $1 - \frac{1}{k+1}$. Indeed in time $f(1 - \frac{1}{k+1})n^{\mathcal{O}(1)} = g(k)n^{\mathcal{O}(1)}$, one then obtains an *exact* solution. A *quasi* PTAS (QPTAS) is an approximation scheme with running time $n^{\text{polylog } n}$, for every $\varepsilon > 0$. Less standardly, we call *subexponential* AS (SUBEXPAS) an approximation scheme with running time 2^{n^γ} for some $\gamma < 1$, for every $\varepsilon > 0$. These approximation schemes can come deterministic or randomised. A maximisation problem Π is *APX-hard* if there is a constant $\gamma < 1$ such that γ-approximating Π is NP-hard. Unless $P = NP$, an APX-hard problem cannot admit a PTAS. Ruling out a SUBEXPAS (under admittedly a stronger assumption than P≠NP) constitutes a sharper inapproximability than an APX-hardness proof. For a minimisation problem, the definitions are the same except that the minus signs are replaced by plus signs and a minimisation problem Π is APX-hard if there is a constant $\gamma > 1$ such that γ-approximating Π is NP-hard.

1.2 Contributions

1.2.1 Clique problems

A very elegant algorithm by Clark, Colbourn and Johnson outputs in polynomial time a maximum clique in a unit disk graph, given with a representation [19]. Finding the complexity of computing a maximum clique in general disk graphs (with arbitrary radii) is a longstanding open problem. However in 2017, Bonnet *et al.*, found a subexponential algorithm and a quasi polynomial time approximation scheme (QPTAS) for maximum clique in disk graphs [9]. The following year, Bonamy *et al.* extended the result to unit ball graphs, and gave a randomised EPTAS for both settings [8]. The current state-of-the-art about the complexity of computing a maximum clique in d-ball graphs is summarised in Table 1.1.

In order to find the complexity of maximum clique in disk graphs, it may be helpful to find superclasses of unit disk graphs for which maximum clique can be solved in polynomial time. Indeed, one

	unit d-ball graphs	general d-ball graphs
$d = 1$	polynomial [39]	polynomial [39]
$d = 2$	polynomial [19]	Unknown but EPTAS [9, 8]
$d = 3$	Unknown but EPTAS [8]	NP-hard [8]
$d = 4$	NP-hard [8]	NP-hard [8]

Table 1.1: Complexity of computing a maximum clique in d-ball graphs

might gain ideas on how to solve the problem in polynomial time in disk graphs, or on the contrary pinpoint the difference between unit disks and disks that makes the problem easy only in unit disk graphs. Similarly, we are interested in extending the EPTAS of Bonamy *et al.* to superclasses of disk graphs. Finally, showing NP-hardness of maximum clique for related graph classes can be helpful in gaining ideas for a NP-hardness proof in disk graphs.

Bonamy *et al.* show that the existence of an EPTAS is implied by the following fact: For any graph G that is a disk graph or a unit ball graph, the disjoint union of two odd cycles is a forbidden induced subgraph in the complement of G. Surprisingly, the proofs for disk graphs on one hand and unit ball graphs on the other hand are not related. Bonamy *et al.* ask whether there is a natural explanation of this common property. They say that such an explanation could be to show the existence of a geometric superclass of disk graphs and unit ball graphs, for which there exists an EPTAS for solving maximum clique.

By looking at Table 1.1, a pattern seems to emerge: The complexity of computing a maximum clique in $(d-1)$-ball graphs and unit d-ball graphs might be related. We extend the question of Bonamy *et al.* and ask for a class of geometric intersection graphs that 1) contains all interval graphs and all unit disk graphs, and 2) for which maximum clique can be solved in polynomial time. Observe that finding such a class also suits our goal of finding superclasses of unit disk graphs for which maximum clique is in P.

Our results on this topic are based on papers that have been published under the following titles:

- *Maximum Clique in Disk-Like Intersection Graphs*, by Édouard Bonnet, Nicolas Grelier and Tillmann Miltzow.
- *Computing a maximum clique in geometric superclasses of disk graphs*, by Nicolas Grelier.

We show that maximum clique can be solved in polynomial time in translates of a fixed bounded convex set, and that there is an EPTAS in homothets of a fixed centrally symmetric bounded convex. Furthermore, our main contribution in [10] is a proof that maximum clique is NP-hard and even APX-hard in intersection graphs of unit disks and axis-parallel rectangles. We provide a new technique to show APX-hardness of maximum clique in intersection graphs for which gadgets are arguably easier to create. We also give partial results towards obtaining an EPTAS for maximum clique in convex pseudo-disk graphs [34].

Our main contribution in [34] is the definition of a class C of geometric intersection graphs which contains all interval graphs and all unit disk graphs, with a proof that maximum clique can be solved efficiently in C. Furthermore, the definition of our class generalises to any dimension, i.e. for any fixed $d \geq 2$ we give a class of geometric intersection graphs that contains all $(d-1)$-ball graphs and all unit d-ball graphs. We conjecture that for $d = 3$, there exists an EPTAS for computing a maximum clique in the corresponding class. It is necessary that these superclasses be defined as classes of geometric intersection graphs. Indeed, it must be if we want to understand better the reason why efficient algorithms exist for both settings. For instance, taking the union of interval graphs and unit disk graphs would not give any insight, since it is a priori not defined

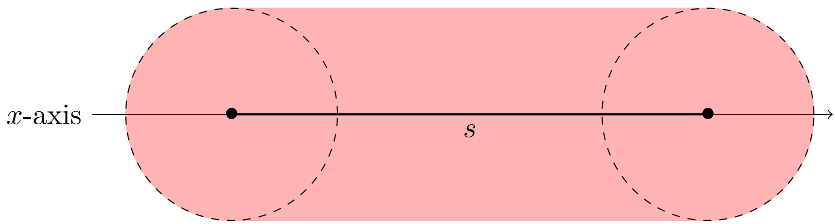

Figure 1.2: The union of the unit disks centred at points of s is a 2-pancake.

by intersection graphs of some geometric objects.

In order to define the class, we first introduce the concept of *d-pancakes*. A 2-pancake is defined as the union of all unit disks whose centres lie on a line segment s on the x-axis. An example is depicted in Figure 1.2. This definition is equivalent to the Minkowski sum of a unit disk centred at the origin and a line segment on the x-axis, where the Minkowski sum of two sets A, B is defined as the set $\{a + b \mid a \in A, b \in B\}$. Similarly a 3-pancake is the union of all unit balls whose centres lie on a disk $\mathcal{D}$ on the xy-plane. More generally, we have:

Definition 1.2. A *d-pancake* is a d-dimensional geometric object. Let us denote by $\{\xi_1, \xi_2, \ldots, \xi_d\}$ the canonical basis of $\mathbb{R}^d$. A d-pancake is defined as the Minkowski sum of the unit d-ball centred at the origin and a $(d-1)$-ball in the linear space spanned by $\{\xi_1, \xi_2, \ldots, \xi_{d-1}\}$.

We denote by Π^d the class of intersection graphs of some finite collection of d-pancakes and unit d-balls. In this paper, we give a polynomial time algorithm for solving maximum clique in Π^2: the intersection graph class of unit disks and 2-pancakes. This is to put in contrast with the fact that computing a maximum clique in intersection graphs of unit disks and axis-parallel rectangles (instead of 2-pancakes) is APX-hard [10], even though maximum clique can be solved in polynomial time in axis-parallel rectangle graphs [43].

Our polynomial time algorithm for maximum clique in Π^2 gives some insight to why the complexity of maximum clique in disk graphs is still unknown. The class of interval graphs is arguably small: there is no induced cycle of length at least 4. Likewise, one can say that the class of unit disk graphs is small, as there is no star with at least 6 leaves. However, with disks one can realise arbitrarily large induced cycles and stars. One could have wondered whether when looking for a geometric class of graphs, wanting both arbitrarily large induced cycles and stars would force too much complexity. Our results with Π^2 shows that actually, this is not where the difficulty lies. Indeed, one can realise with unit disks and 2-pancakes arbitrarily large induced cycles and stars. To solve maximum clique in disk graphs, or to show NP-hardness, it seems a good idea to investigate what are the disk graphs that are not in Π^2.

1.2.2 Plane geometric graphs on degenerate point sets

We study problems on plane geometric graphs on point sets where several points can lie on a line, i.e. degenerate point sets. There are problems for which the difference between general position and degenerate position is negligible. Let us briefly state an example here: Given a simple polygon and two points p and q, check whether p and q are on the same side of the polygon. To solve the problem, a solution consists in considering the segment $[p, q]$, and then counting how many segments of the polygon it intersects. If the point set P defined as the vertices of the polygon union p and q is in general position, observe that p and q are on the same side if and only if $[p, q]$ intersects an even number of segments of the polygon. If P is in degenerate position, the situation can only be slightly more difficult. Indeed, it might occur that $[p, q]$ intersects a vertex r of the polygon. Then, one has to check whether $[p, q]$ is actually traversing the polygon at r, or simply touching it. Thus, this is an example where allowing for degenerate point set only requires slightly more care.

In this thesis, we study two problems on degenerate point sets, where there is a fundamental difference with the general position setting. The results are based on papers with the following titles:

- *Hardness and approximation of minimum convex partition*, by Nicolas Grelier.
- *Connectivity of bistellar flip graphs of planar point sets in special position*, by Nicolas Grelier and Emo Welzl.

The first paper has appeared in the proceedings of the 38th International Symposium on Computational Geometry (SoCG 2022). The second paper is not submitted yet.

The first problem we present is the one of finding a convex partition of a point set with as few faces as possible. It is known that there is an $\frac{8}{3}$-approximation algorithm running in polynomial time for point sets in general position [52, 71]. However for point sets that can be degenerate no $\mathcal{O}(n^{1-\varepsilon})$-approximation algorithm was known for any $\varepsilon > 0$. Our main contributions are an $\mathcal{O}(\log(OPT))$-approximation algorithm running in polynomial time that works for all point sets, and a proof of NP-hardness [35]. We also present several other approximations algorithms, some with faster running time at the cost of a worse approximation ratio, some algorithms that require additional restrictions on the point set, and an FPT approximation algorithm with respect to the number of faces in a minimum solution.

The second problem concerns the connectivity of the bistellar flip graph of triangulations on a point set. In this graph, there is a vertex for each triangulation, and two triangulations are adjacent if they are sufficiently similar, as characterised by the notion of flip: a minimal change to a triangulation that yields another triangulation. It was shown by Wagner and Welzl that the bistellar flip graph is $(n-3)$-connected for point sets in the plane in general position [82]. We extend this result to point sets in degenerate position. Our main contribution is the finding of a point set in degenerate position for which one of the lemmas in [82] does not hold, with a proof on how to circumvent this issue for extending the main theorem to the degenerate setting.

1.2.3 Other contributions

We give here an exhaustive list of the papers on which we worked during the doctoral studies, but which are not part of this thesis. It includes two papers about colourings, one paper about the VC-dimension of convex sets, and one about hyperplane transversals in high dimensions.

- *Approximate Strong Edge-Colouring of Unit Disk Graphs* [37], by Nicolas Grelier, Rémi de Joannis de Verclos, Ross J. Kang and François Pirot.
- *Nearest-Neighbor Decompositions of Drawings* [20], by Jonas Cleve, Nicolas Grelier, Kristin Knorr, Maarten Löffler, Wolfgang Mul-zer and Daniel Perz.
- *On the VC-dimension of convex sets and half-spaces* [36], by Nicolas Grelier, Saeed Gh. Ilchi, Tillmann Miltzow and Shakhar Smorodinsky.
- *Well-Separation and Hyperplane Transversals in High Dimensions* [7], by Helena Bergold, Daniel Bertschinger, Nicolas Grelier, Wolfgang Mulzer and Patrick Schnider.

CHAPTER 2

Maximum clique in superclasses of unit disk graphs

2.1 Introduction

The study of unit disk graphs really started in 1990 after the publication of the suitably titled paper “Unit disk graphs” by Clark, Colbourn and Johnson [19]. Along with unit disk graphs, they consider the class of grid graphs. A graph is a *grid graph* if there is an injective map from its vertex set to points in the plane with integer coordinates, such that two vertices are adjacent if and only if their L_1 distance is equal to 1. It is easy to observe that grid graphs are a subclass of unit disk graphs. In 1990, it was already known that the dominating set problem, hamiltonicity or the Steiner tree problem are NP-hard in grid graphs [32, 46, 48]. This immediately implies that the same problems are also NP-hard in unit disk graphs. Moreover, Clark, Colbourn and Johnson show that 3-colouring and maximum independent set in unit disk graphs are NP-hard too [19]. It seems that most problems that are NP-hard for general graphs are also NP-hard for unit disk graphs. Surprisingly, they show that maximum clique is an exception when a unit disk representation is given.

Theorem 2.1 ([19]). *Given a representation of a unit disk graph with n vertices, one can compute a maximum clique in $\mathcal{O}(n^{4.5})$-time.*

Let us state here a sketch of the proof. First, we define the notions of *lens* and *half-lens*.

Definition 2.2 (Lens and half-lens). We denote by $\mathcal{D}(c, \rho)$ a closed disk centred at c with radius ρ. Let $\mathcal{D} = \mathcal{D}(c, \rho)$ and $\mathcal{D}' = \mathcal{D}(c', \rho')$ be two intersecting disks. We call *lens induced by $\mathcal{D}$ and $\mathcal{D}'$* the region $\mathcal{D} \cap \mathcal{D}'$. We call *half-lenses* the two closed regions obtained by dividing the lens along the line (c, c').

Observe that in Definition 2.2, we consider disks of any radius instead of unit disks. In the proof of Theorem 2.1, the authors first observe that the diameter of a half-lens of two disks with radius r whose centres are at distance r is equal to r. Therefore, if r is at most 2 then two unit disks whose centres are in such a half-lens intersect. Now to compute a maximum clique in a unit disk graph given with a representation, first guess in $\mathcal{O}(n^2)$ two most remote centres c and c' of unit disks in a maximum clique. Let us denote by ρ the distance between c and c'. By assumption, all the centres of the other unit disks in that maximum clique are in the lens induced by $\mathcal{D}(c, \rho)$ and $\mathcal{D}(c', \rho)$. To compute the maximum clique, one has simply to compute a maximum clique among the unit disks whose centres are in that lens. The set of unit disks whose centres are in the same half-lens induce a clique, because the diameter of the half-lens is equal to ρ, which is at most 2. Therefore the problem boils down to computing a maximum clique in a cobipartite graph, which can be done in $\mathcal{O}(n^{2.5})$ [28].

A question follows naturally: What if the unit disk graph is given as an abstract graph, without a representation? Is it possible to find such a representation, and then use the algorithm by Clark, Colbourn and Johnson [19]? This cannot be done in polynomial time, unless $P = NP$, as shown by Breu and Kirkpatrick [13]. Evenmore, Kang and Müller showed that recognition of unit disk graphs is $\exists\mathbb{R}$-hard [49].

However, Raghavan and Spinrad [70] showed in 2003 the existence of an algorithm running in polynomial time that given an abstract graph G, either outputs a maximum clique of G or a certificate that G is not a unit disk graph. Note that even if G is not a unit disk graph, the algorithm may still output a maximum clique. Such an algorithm which does not require a representation is called *robust*. They introduce the concept of *Cobipartite Neighbourhood Edge Elimination Ordering* (CNEEO). Let G be a graph with m edges. Let $\Lambda = e_1, e_2, \ldots, e_m$ be an ordering of the edges. Let $G_{\Lambda,k}$ be the subgraph of G with edge set $\{e_k, e_{k+1}, \ldots, e_m\}$. For each $e_k = \{u, v\}$, $N_{\Lambda,k}$ is defined as the set of vertices adjacent to u and v in $G_{\Lambda,k}$.

Definition 2.3 ([70])**.** An edge ordering $\Lambda = \{e_1, e_2, \ldots, e_m\}$ is a CNEEO if for each e_k, $N_{\Lambda,k}$ induces a cobipartite graph in G.

Lemma 2.4 ([70])**.** *Given a graph G and a CNEEO Λ for G, a maximum clique in G can be found in $\mathcal{O}(n^{4.5})$-time.*

The algorithm consists in computing a maximum clique in each $N_{\Lambda,k}$, for $1 \leq k \leq m$. As $N_{\Lambda,k}$ is a cobipartite graph, the total running time is in $\mathcal{O}(n^{4.5})$ [28].

They propose a greedy algorithm for finding a CNEEO: When having chosen the first $i-1$ edges $e_1, \ldots, e_{i-1}$, try every remaining edge one by one until finding one e_i that satisfies the required property. If no such edge exists, return that the graph does not admit a CNEEO, which follows from Lemma 2.5.

Lemma 2.5 ([70])**.** *If G admits a CNEEO, then the greedy algorithm finds a CNEEO for G.*

Assume the first $i-1$ edges are chosen by this greedy algorithm. There are $\mathcal{O}(n^2)$ remaining edges to consider, and it takes $\mathcal{O}(n^{2.5})$-time to check whether a new edge satisfies the conditions. Therefore, the total running time to check if a graph has a CNEEO, and to output it if it has one, is in $\mathcal{O}(n^{6.5})$. Raghavan and Spinrad show that all unit disk graphs have a CNEEO (by considering edges in decreasing order of length in any fixed representation), which implies that one can compute a maximum clique in unit disk graphs, even without a representation, in time $\mathcal{O}(n^{6.5})$.

In this section, we show that the class of intersection graphs of translates of a fixed convex set, as well as the class Π^2, are subclasses of the class of graphs which admit a CNEEO. It immediately follows that one can compute a maximum clique in those classes in time $\mathcal{O}(n^{6.5})$.

2.2 Translates of a fixed set

We show in this section that we can extend the algorithm of Clark, Colbourn and Johnson [19] and its robust version [70] from unit disks to any centrally symmetric, bounded, convex set.

Theorem 2.6. *Maximum clique admits a robust algorithm in running in $\mathcal{O}(n^{6.5})$-time in intersection graphs of translates of a fixed centrally symmetric, bounded, convex set.*

Moreover, as shown by Aamand et al. [1], for every bounded and convex set S_1, there exists a centrally symmetric, bounded and convex set S_2 such that $\mathcal{G}_{S_1} = \mathcal{G}_{S_2}$, where $\mathcal{G}_S$ denotes the intersection graphs class of translates of S. Thus we obtain the immediate corollary:

Corollary 2.7. *Maximum clique admits a robust algorithm in running in $\mathcal{O}(n^{6.5})$-time in intersection graphs of translates of a fixed bounded and convex set.*

We prove Theorem 2.6 in two steps. First we show how to compute in polynomial time a maximum clique when a representation is given. Secondly we use the result by Raghavan and Spinrad [70] to obtain a robust algorithm.

2.2.1 Maximum clique with a representation

Let S be a centrally symmetric, bounded, convex set. We can define a corresponding norm as follow: for any $x \in \mathbb{R}^2$, let $\|x\|$ be equal to $\inf\{\lambda > 0 \mid x \in \lambda S\}$. This is well-defined since S is bounded. It is absolutely homogeneous because S is centrally symmetric, and it is subadditive because S is convex. Therefore $\|.\|$ is a norm. Let S_1 and S_2 be two translates of S, with respective centres c_1 and c_2. Observe that S_1 and S_2 intersect if and only if $\|c_1 - c_2\| \leqslant 2$. Let us assume that $d := \|c_1 - c_2\| \leqslant 2$. We denote by S' the set S scaled by d: $S' := dS$, and we then define: $D := \{x \in \mathbb{R}^2 \mid \|x - c_1\| \leqslant d, \|x - c_2\| \leqslant d\}$. Equivalently we have $D = (c_1 + S') \cap (c_2 + S')$. Figure 2.1 shows an example. If S is a unit disk, then D is the lens induced by two disks with radius d, such that the boundary of one contains the centre of the other.

Lemma 2.8. *The set D is centrally symmetric around $c := (c_1 + c_2)/2$.*

Proof. Let x be a point in D, we need to show that $\bar{x} := x + 2(c - x)$ is in D too. As $D = (c_1 + S') \cap (c_2 + S')$, it is sufficient to show $\bar{x} \in c_1 + S'$ and $\bar{x} \in c_2 + S'$. By definition, $\bar{x}$ is equal to $c_1 + c_2 - x$. Since x is in D, then $\|c_2 - x\| \leqslant d$, which implies that $c_2 - x$ is in

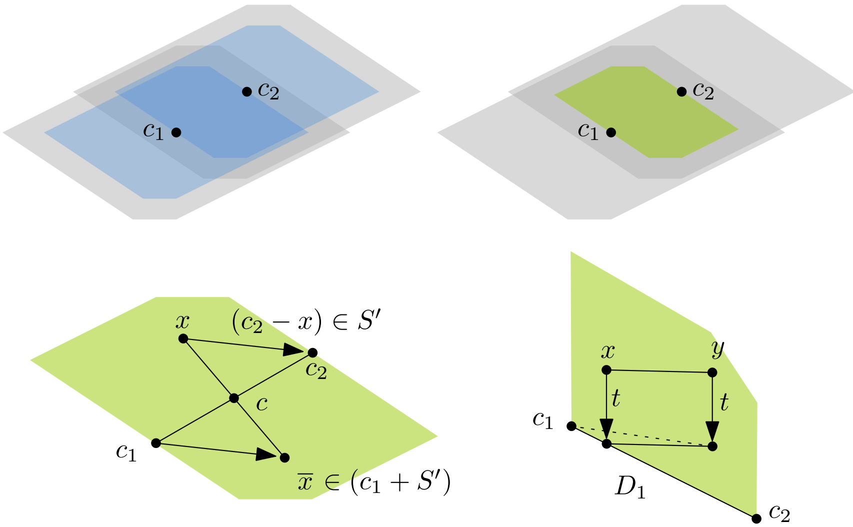

Figure 2.1: Top left: The grey sets are scaled about their centre so that the centre of one set is on the boundary of the other. Top right: the intersection D. Bottom left: Illustration of Lemma 2.8. Bottom right: Illustration of Lemma 2.10.

S'. Therefore $\bar{x}$ is in $c_1 + S'$. By the symmetry of the arguments, we obtain that $\bar{x}$ is in D. □

Lemma 2.9. *There exist parallel tangents to D at c_1 and c_2.*

Proof. Let us denote by ℓ_1 a tangent to D at c_1. Then we denote by ℓ_2 the line parallel to ℓ_1 that contains c_2. We claim that ℓ_2 is tangent to D. By construction D is convex, as the intersection of two convex sets. This implies that ℓ_2 is tangent to D if and only if $D \cap \ell_2$ is a line segment that contains c_2. This line segment may be only one point. Let x be a point in $D \cap \ell_2$. By Lemma 2.8, D is centrally symmetric around c. Therefore $x + 2(c - x)$ is in D, and by construction it is also in ℓ_1. Since $D \cap \ell_1$ is a line segment that contains c_1, thus $D \cap \ell_2$ is a line segment that contains c_2. □

We cut D along the line ℓ going through c_1 and c_2, and split D into two sets denoted by D_1 and D_2. We define D_1 as the set of points below this line, and D_2 as the set of points not below. We have the following lemma:

Lemma 2.10. *Let i be in $\{1, 2\}$, and let x and y be in D_i. Then we have $\|x - y\| \leqslant d$.*

Proof. We do the proof for $i = 1$, and the case $i = 2$ can be done symmetrically. By Lemma 2.9, there are parallel tangents ℓ_1 and ℓ_2 of D at c_1 and c_2. Without loss of generality, let us assume that they are vertical, that c_1 is to the left of c_2 and x to the left of y. We denote by $\tilde{x}$ (respectively $\tilde{y}$) the vertical projection of x (respectively y) on ℓ. Without loss of generality $\|x - \tilde{x}\| \leqslant \|y - \tilde{y}\|$. We define $t = x - \tilde{x}$. Note that $\|x - y\| = \|(x - t) - (y - t)\| = \|\tilde{x} - (y - t)\|$. Furthermore, we can move $\tilde{x}$ on ℓ towards c_1 and this will only increase the distance to $(y - t)$. We get $\|\tilde{x} - (y - t)\| \leqslant \|c_1 - (y - t)\|$. By definition $(y - t) \in D_1 \subset D$ and thus $\|c_1 - (y - t)\| \leqslant d$. This implies $\|x - y\| \leqslant d$ and finishes the proof. □

Following the arguments of Clark, Colbourn and Johnson [19], one first guesses in quadratic time S_1 and S_2 in a maximum clique C such that the distance between their centres $\|c_1 - c_2\|$ is maximised among the pairs $S_1, S_2 \in C$. One can then remove all the objects not centred in D. By Lemma 2.10, the intersection graph induced by the sets centred in D is cobipartite. Since computing an independent set in a bipartite graph can be done in $\mathcal{O}(n^{2.5})$-time, one can compute a maximum clique in G in $\mathcal{O}(n^{4.5})$-time.

2.2.2 Maximum clique without a representation

Proof of Theorem 2.6. By what was argued in Section 2.1, it is sufficient to show that for any centrally symmetric, bounded, convex set S, and any intersection graph G of translated of S, there exists a CNEEO on G. Let us consider such a graph G with a representation. Arguing with Lemma 2.10 as previously, ordering the edges by non-increasing length gives a CNEEO, where the length of an edge is the distance between the two centres. □

2.3 Axis-parallel rectangles and unit disks

In Sections 2.3 and 2.4, we consider intersection graphs of two types of geometrical objects. In this section, we consider intersection graphs of axis-parallel rectangles and unit disks, whereas in Section 2.4, we consider intersection graphs of 2-pancakes and unit disks. Recall that we denote this class by Π^2. In Section 2.4, we show how to compute a maximum clique in Π^2 in polynomial time. To motivate better this result, we show here that maximum clique is NP-hard, and even APX-hard, for intersection graphs of axis-parallel rectangles and unit disks. This can be surprising, as maximum clique in axis-parallel rectangles is in P [43], and maximum clique in unit disk graphs is in P too by Theorem 2.1.

2.3.1 The co-2-subdivision approach

Let us first discuss a technique, the *co-2-subdivision approach*, that has been used to show hardness of maximum clique in intersection graphs. Maximum independent set, which boils down to maximum clique in the complement graphs, is APX-hard on subcubic graphs [3]. A folklore self-reduction first discovered by Poljak [68] consists of subdividing each edge of the input graph twice (or any even number of times). One can show that this reduction preserves the APX-hardness. Therefore, a way to establish such an intractability for maximum clique on a given intersection graph class is to show that for every (subcubic) graph G, its complement of 2-subdivision $\overline{\text{Subd}_2(G)}$ (or $\overline{\text{Subd}_s(G)}$ for a larger even integer s, see [30]) is representable. Maximum independent set admits a PTAS on planar graphs, but remains NP-hard. Hence showing that for every (subcubic) planar graph G, the complement of an even subdivision of G is representable shows the simple NP-hardness (see [16, 30]).

So far, representing complements of even subdivisions of graphs belonging to a class on which maximum independent set is NP-hard (respectively APX-hard) has been the main, if not unique[1], approach to show the NP-hardness (respectively APX-hardness) of maximum clique in geometric intersection graph classes. This approach was used by Middendorf and Pfeiffer [62] for some restriction of string graphs, the so-called B_1-*VPG graphs*, by Cabello et al. [16] to settle the then long-standing open question of the complexity of maximum clique for segments (with the class of planar graphs), by Francis et al. [30] for 2-interval, unit 3-interval, 3-track, and unit 4-track graphs (with the class of all graphs; showing APX-hardness), and unit 2-interval and unit 3-track graphs (with the class of planar graphs; showing only NP-hardness), by Bonnet et al. [9] for filled ellipses and filled triangles, and by Bonamy et al. [8] for ball graphs, and 4-dimensional unit ball graphs. Bonnet et al. [9] show that the complement of two mutually induced odd cycles is not a disk graph. As a consequence, to show the NP-hardness of maximum clique on disk graphs with the described approach, one can only hope to represent all the graphs without two mutually induced odd cycles. However we do not know if maximum independent set is even NP-hard in that class.

The main conceptual contribution of this section is to suggest an

[1] Admittedly Butman et al. [15] showed that maximum clique is NP-hard on 3-interval graphs, by reducing from Max 2-DNF-SAT which is very close to Max Cut. However this result was later subsumed by [30].

alternative to that approach. We introduce a technical intermediate problem that we call Max Interval Permutation Avoidance (MIPA, for short), which is a convenient way of seeing Max Cut. We prove that MIPA is unlikely to have an approximation scheme running in subexponential time. We then transfer that lower bound to maximum clique in the intersection graphs of objects that can be either unit disks or axis-parallel rectangles; a class for which the *co-2-subdivision* approach does not seem to work. Recall that when all the objects are unit disks or when all the objects are axis-parallel rectangles, polynomial-time algorithms are known.

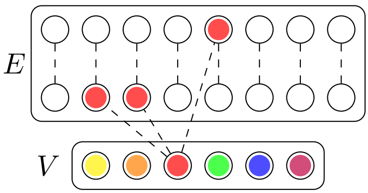

(a) Co-2-subdivision of subcubic graphs: edges are represented by an antimatching, vertices, by a clique.

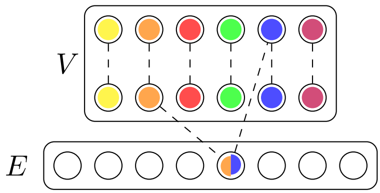

(b) MIPA approach: vertices are represented by an antimatching with constant weight, edges, by a clique.

Figure 2.2: Dashed segments represent non-edges. Both the *co-2-subdivision* and the MIPA approaches require to construct an antimatching and a clique. In the *co-2-subdivision* approach, the *clique vertices* have co-degree 3 to the antimatching. In the MIPA approach their co-degree is only 2. While the difference is seemingly small, the graph class formed by axis-parallel rectangles and unit disks is not amenable to the *co-2-subdivision* approach (see Section 2.3.3).

Anticipating on Section 2.3.3 where MIPA is defined, one can already see on Figure 2.2 that both approaches require to represent an antimatching (i.e., a complement of an induced matching), a clique, and some relation between them. Antimatchings (and obviously cliques) of arbitrary size are representable by half-planes and unit disks. The difficulty in both cases is to get the right adjacencies between the antimatching and the clique. The MIPA approach only needs the vertices of the clique to *avoid* two vertices in the antimatching, whereas this number is at least three in the *co-2-subdivision* approach. This seemingly small difference is actually crucial, as we will see in Section 2.3.3.

2.3.2 Some background inapproximability results

The Exponential-Time Hypothesis (ETH, for short) of Impagliazzo and Paturi [44] asserts that there is no subexponential-time algorithm solving k-SAT. More precisely, for every integer $k \geqslant 3$, there is an $\varepsilon > 0$ such that k-SAT cannot be solved in time $2^{\varepsilon n}$ on n-variable instances. If we define s_3 (taking the same notation as in the original paper) as the infimum of the reals δ such that 3-SAT can be solved in time $2^{\delta n}$, then the ETH can be expressed as $s_3 > 0$. Impagliazzo *et al.* [45] present a subexponential-time Turing-reduction parameterized by a positive real $\varepsilon > 0$ which, given a k-SAT-instance ϕ with n variables and m clauses, produces at most $2^{\varepsilon n}$ k-SAT-instances $\phi_1, \ldots, \phi_t$ such that $\phi \Leftrightarrow \bigvee_{i \in [t]} \phi_i$, each ϕ_i having no more than n variables and $C_\varepsilon n$ clauses for some constant C_ε (depending solely on ε, and *not* on n and m). This important reduction is known as the Sparsification Lemma. One can observe that, due to the Sparsification Lemma, there is an $\varepsilon > 0$ such that there is no algorithm solving k-SAT in time $2^{\varepsilon m}$ on m-clause instances, assuming that the ETH holds. For the sparsification of a 3-SAT-instance, the constant C_ε can be upper-bounded by $10^8(1/\varepsilon)^2 \log^2(1/\varepsilon)$. One can sparsify a 3-SAT-instance in $2^{\frac{s_3}{2}n}$ instances with at most $X := C_{s_3/2} \leqslant 10^8(2/s_3)^2 \log^2(2/s_3)$ occurrences per variable. Assuming the ETH, these sparse instances cannot be solved in time $2^{\frac{s_3}{2}n}$.

Let us state some inapproximability results, that we use for our reduction. A detailed proof of these theorems can be found in [10].

Theorem 2.11. *[41, 64, 45] Under the ETH, for every* $\delta > 0$ *one cannot distinguish in time* $2^{n^{1-\delta}}$, *n-variable m-clause* 3-SAT-*instances that are satisfiable from instances where at most* $(7/8 + o(1))m$ *clauses can be satisfied, even when each variable appears in at most X clauses. Thus* 3-SAT-X *cannot be* $7/8 + o(1)$*-approximated in time* $2^{n^{1-\delta}}$.

Now, we recall the definition of Not-All-Equal k-SAT (NAE k-SAT, for short).

> Not-All-Equal k-SAT
> **Input:** A conjunction of m "clauses" $\phi = \bigwedge_{i\in[m]} C_i$ each on at most k literals.
> **Goal:** Find a truth assignment of the n variables such that each "clause" has at least one satisfied literal and at least one non-satisfied literal.

The Not-All-Equal k-SAT-B-problem is the same but each variable appears in at most B clauses (similarly as for k-SAT-B). The adjective *Positive* prepended to a satisfiability problem means that no negation (or *negative literal*) can appear in its instances. As a slight abuse of notation, we keep the same problem names for the maximisation versions, where all the clauses may not be simultaneously satisfied but the goal is to satisfy the largest fraction of them. Another abuse of notation is that we call *clauses* the *not-all-equal* constraints, and still denote them with $\vee$. The performance guarantee of an approximation algorithm is then defined as the minimum of *number of satisfied clauses/m* taken over all the instances.

We recall that X is a function of the value s_3 assumed to be positive by the ETH.

Theorem 2.12 ([10])**.** *Under the ETH, for every* $\delta > 0$ *one cannot distinguish in time* $2^{n^{1-\delta}}$, *n-variable m-clause Positive Not-All-Equal*

3-SAT-3-instances that are satisfiable from instances where at most γm clauses can be satisfied, with $\gamma := (60000X^2 - 9)/(60000X^2)$. Thus Positive Not-All-Equal 3-SAT-3 cannot be γ-approximated in time $2^{n^{1-\delta}}$.

Corollary 2.13 ([10]). *Approximating* Positive NAE 3-SAT-3 *within ratio* 49888956/49888957 *is NP-hard.*

2.3.3 MIPA, unit disks and rectangles

We first introduce the Max Interval Permutation Avoidance-problem (MIPA, for short), a convenient intermediate problem to show APX-hardness for geometric problems. We start with an informal description. Let M be a perfect matching between the n points $[n] \times \{0\}$ and $[n] \times \{1\}$, in $\mathbb{N}^2$. This matching can be represented by a permutation σ, such that for every $i \in [n]$, $(i, 0)$ is matched with $(\sigma(i), 1)$. Imagine now a set of intervals on the line $y = 1/2$ whose endpoints are all in $[n]$. The aim is to move each interval "up" or "down", by translating it by $(0, 1/2)$ or by $(0, -1/2)$, respectively, such that the number of edges of M with no endpoint on a translated interval is maximised. An edge of M with at least one endpoint in a *moved* (or *positioned*) interval is said *covered* or *destroyed*. The edge is said *uncovered* or *preserved* otherwise. Equivalently Max Interval Permutation Avoidance aims to minimising the number of covered edges, or maximising the number of uncovered edges. We choose the maximisation formulation, since we will both reduce from a maximisation problem (Positive Not-All-Equal 3-SAT-3) and to a maximisation problem (maximum clique in disks and rectangles). Thus the objective value will be the number of *uncovered edges*.

Max Interval Permutation Avoidance
Input: A permutation σ over $[n]$ representing a perfect matching M between the points $(1,0),(2,0),\ldots,(n,0)$ and $(\sigma(1),1),(\sigma(2),1),\ldots,(\sigma(n),1)$ respectively, and a set of integer ranges $\mathcal{I} := \{I_1,\ldots,I_h\}$ where $I_k := [\ell_k, r_k]$ and $1 \leqslant \ell_k \leqslant r_k \leqslant n$, for every $k \in [h]$.
Goal: A placement function $p : \mathcal{I} \to \{0,1\}$ encoding that interval I_k has its endpoints positioned in $(\ell_k, p(I_k))$ and $(r_k, p(I_k))$, which maximises the number of edges of M with no endpoint on a positioned interval.

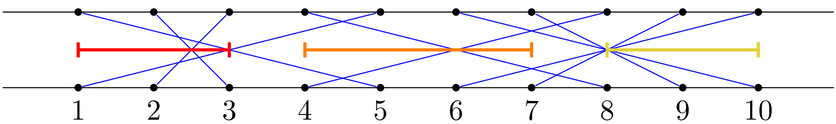

Figure 2.3: An example of a symmetric instance of MIPA with three disjoint ranges.

A MIPA-input may interchangeably be given as $(\sigma,\mathcal{I})$ or as $(M,\mathcal{I})$. One may observe that a *constant* placement (i.e., $p(I_1) = \ldots = p(I_h) = 0$, or $p(I_1) = \ldots = p(I_h) = 1$) is a worse solution when the intervals of $\mathcal{I}$ span $[n]$, since it covers all the edges of M. We say that the matching M is *symmetric* if $(i,0)(j,1) \in M$ implies that $(i,1)(j,0) \in M$, for every $i,j \in [n]$; in the geometric viewpoint, it is equivalent to $y = 1/2$ being a symmetry axis of M, and in the permutation viewpoint, it is equivalent to σ being a product of pairwise-disjoint transpositions. Other handy (as far as hardness of geometric problems is concerned) technical problems involving intervals and/or permutations include crossing-avoiding matching in Guśpiel [40] or crossing-minimising perfect matching in Guśpiel et al. [2], the problem of covering a 2-track point set by selecting k 2-track intervals [61] or structured 2-track hitting set [11]. It is no coincidence that these convenient starting problems all involve matchings/permutations and/or intervals. Indeed the latter objects are more easily encoded in a geometric setting than their generalizations: arbitrary binary relations and arbitrary sets. Later we will see how disks can encode intervals and how rectangles can encode a permutation, in the context of the maximum clique problem.

We rule out an approximation scheme for Max Interval Permutation Avoidance, even if subexponential-time is allowed. In particular a QPTAS for MIPA is highly unlikely. We recall that $\gamma = (60000X^2 - 9)/(60000X^2)$ and that X is a finite integral constant, assuming the ETH ($s_3 > 0$).

Lemma 2.14. *For every $\delta > 0$, Max Interval Permutation Avoidance cannot be γ'-approximated in time $2^{|M|^{1-\delta}}$, with $\gamma' := 1 - (1 - \gamma)/13 < 1$, unless the ETH fails. Besides Max Interval Permutation Avoidance is NP-hard and APX-hard. These results hold even if the length of every interval of $\mathcal{I}$ is at most 5, and the matching M is symmetric.*

Proof. We give a reduction ϕ from Positive Not-All-Equal 3-SAT-3 to Max Interval Permutation Avoidance. Let ϕ be a Positive NAE 3-SAT-3-instance, with variables $x_1, \ldots, x_n$ and clause $C_1, \ldots, C_m$. For every $x_i \in C_j$, we denote by $\text{occ}(x_i, C_j)$ the number of occurrences of x_i in the clauses $C_1, \ldots, C_j$. We observe that $\text{occ}(x_i, C_j) \in \{1, 2, 3\}$. We build an instance $\rho(\phi) := (M, \mathcal{I})$ of MIPA in the following way. For each variable x_i of ϕ, we reserve a range $[3(i-1)+1, 3(i-1)+3]$ with 3 integral points on both lines $y = 0$ and $y = 1$. These points will be matched by M to points in the clause gadgets. We add the interval $X_i := [3(i-1)+1, 3(i-1)+3]$ to $\mathcal{I}$. We now describe the 2-clause and the 3-clause gadgets.

For every 2-clause $C_j := x_a \vee x_b$, we allocate a slot S_j of size 9 (on $y = 0$ and $y = 1$) appended to the current last position. The first half of S_j, that is, the indices in $[s_j, s_j + 4]$ of S_j correspond to x_a, while the indices in $[s_j + 5, s_j + 9]$ correspond to x_b. For every $(d_1, d_2) \in \{(0,1),(1,0)\}$ and $h \in [4]$, we add to M the edge between $(s_j + h, d_1)$ and $(s_j + 5 + h, d_2)$. We add the intervals $C_j(x_a) := [s_j, s_j + 4]$ and $C_j(x_b) := [s_j + 5, s_j + 9]$ to $\mathcal{I}$. Finally for each $(d_1, d_2) \in \{(0,1),(1,0)\}$, we add to M the edges between (s_j, d_1) and $(3(a-1)+\text{occ}(x_a, C_j), d_2)$, and between (s_j+3, d_1) and $(3(b-1) + \text{occ}(x_b, C_j), d_2)$.

For every 3-clause $C_j := x_a \vee x_b \vee x_c$, we allocate a slot S_j of size 15 (on $y = 0$ and $y = 1$) appended to the current last position. The first third of S_j, that is, the indices in $[s_j, s_j+4]$ of S_j correspond to x_a, the second third, the indices in $[s_j + 5, s_j + 9]$ correspond to x_b, and the last third, the indices in $[s_j + 10, s_j + 14]$ correspond to x_c. We add the intervals $C_j(x_a) := [s_j, s_j+4]$, $C_j(x_b) := [s_j+5, s_j+9]$, and $C_j(x_c) := [s_j + 10, s_j + 14]$ to $\mathcal{I}$. Similarly for every $(d_1, d_2) \in \{(0,1),(1,0)\}$ and $(h, p) \in \{(a,0),(b,1),(c,2)\}$, we add to M the edge between $(s_j + 5p, d_1)$ and $(3(h-1) + \text{occ}(x_h, C_j), d_2)$. We call these edges *internal* (same for the 2-clause gadget). Finally we add to M four edges from every pair of ranges in $\{[s_j, s_j+4], [s_j+5, s_j+9], [s_j+10, s_j+14]\}$, two starting on the line $y = 0$ (ending on $y = 1$) and two starting on $y = 1$ (ending on $y = 0$). We call these edges *variable-clause* (same for the 2-clause gadget).

For each variable x_i with only two occurrences in ϕ, we link its third occurrence pair to a dummy pair $(d_i, 0), (d_i, 1)$, appended to the current last position. That is, we add the edges $(3(i-1)+3, 0)(d_i, 1)$ and $(3(i-1)+3, 1)(d_i, 0)$ to M. Although not needed, we also add the singleton interval $D_i := \{d_i\}$ to $\mathcal{I}$. We call it *dummy gadget* and consider it as a special case of a clause gadget. This finishes the construction of the MIPA-instance $(M, \mathcal{I})$. Observe that every point is matched, and that all the intervals of $\mathcal{I}$ are pairwise disjoint, and of length at most 5. The perfect matching M comprises at most $3n + 15m + n \leqslant 49n$ edges.

We assume that ϕ is satisfiable, and let $\mathcal{V}$ be a satisfying assignment. We build the following solution to the MIPA-instance. We push the interval X_i up (to the line $y = 1$) if x_i is set to true by $\mathcal{V}$, and we push it down (to the line $y = 0$) otherwise. In the clause gadgets (and dummy gadgets), we do the opposite: we push $C_j(x_i)$ (D_i) down if x_i is set to true, and up if x_i is set to false. This solution preserves four edges within each clause gadget, and an additional $3n$ edges between the variable gadgets and the clause gadgets. Hence the total number of preserved edges is $4m + 3n$.

We now assume that at most γm clauses of ϕ are satisfiable. Let p be a placement function of the intervals of $\mathcal{I}$, maximising the number of preserved edges of M. We first argue that not giving the same placement (up/1 or down/0) to the three (respectively two) intervals $C_j(x_a), C_j(x_b), C_j(x_c)$ (respectively $C_j(x_a), C_j(x_b)$) of a 3-clause gadget (respectively 2-clause gadget) is always better. Note that any equal placement destroys all the edges of M internal to the clause gadget of C_j, and preserves at most three variable-clause edges. On the other hand, a placement with at least one interval on each side preserves already four internal edges. We can then assume that p does not give equal placement in any clause gadget. Let $\mathcal{V}$ be the assignment of the variables of ϕ which sets x_i to true if $p(X_i) = 1$, and to false, if $p(X_i) = 0$. By assumption $\mathcal{V}$ does not satisfy at least $(1 - \gamma)m$ clauses. In each corresponding clause gadget, one can preserve at most two variable-clause edges of M. Indeed all three variable-clause edges incident to the clause gadget and not covered by the placement of the X_i land on the same side. By the previous remark, at least one such edge should be destroyed (to preserve four internal edges). Thus the placement p preserves at most $3n + 4m - (1 - \gamma)m$ edges.

Since $|M| = O(n+m) = O(n)$ and $\frac{3n+4m-(1-\gamma)m}{3n+4m} \leqslant 1 - \frac{1-\gamma}{13}$, by Theorem 2.12 MIPA cannot be γ'-approximated in time $2^{|M|^{1-\delta}}$, under the ETH. Besides, by Corollary 2.13, MIPA cannot be approximated with ratio 648556435/648556436 in polynomial-time, unless P=NP. In particular, this problem is NP-hard and even APX-hard. $\square$

We recall that maximum clique can be solved in polynomial-time in unit disk graphs [19, 70] and in axis-parallel rectangle intersection graphs [14]. Now if the objects can be unit disks *and* axis-parallel rectangles, we show that even a SUBEXPAS is unlikely.

Theorem 2.15. *For every $\delta > 0$, maximum clique in intersection graphs G of unit disks and axis-parallel rectangles cannot be c-approximated in time $2^{|V(G)|^{1-\delta}}$, with $c := 1 - (1-\gamma)/153 < 1$, unless the ETH fails. Besides this problem is NP-hard and APX-hard.*

Proof. We give a reduction from Max Interval Permutation Avoidance to maximum clique in intersection graphs of unit disks and axis-parallel rectangles, that also holds for maximum clique in intersection graphs of half-planes and axis-parallel rectangles. Let $(M, \mathcal{I})$ be an instance of MIPA over $[n]$, where M is symmetric, and all the intervals of $\mathcal{I}$ have size at most 5. We build the following set of axis-parallel rectangles $\mathcal{R}$ and half-planes $\mathcal{H}$. See Figure 2.4 for an illustration.

Let O be the origin of the plane. We place from left to right $n+2$ points $p_0, p_1, \ldots, p_n, p_{n+1}$ on a convex curve in the top-left quadrant, say $x \mapsto -1/x$ on $[-(1+\lambda), -1]$ for some small $\lambda > 0$. We wiggle the points p_i so that for every $i \leqslant j \in [n]$, the slope of the line passing through middle(p_{i-1}, p_i) and middle(p_j, p_{j+1}) has a distinct value. We define $q_0, q_1, \ldots, q_n, q_{n+1}$, such that O is the middle of the segment $p_i q_i$ for every $i \in [0, n+1]$. In other words, this new chain is obtained by central symmetry about O. Observe that sorted by x-coordinates, these $2n+4$ points read $p_0, p_1, \ldots, p_n, p_{n+1}, q_{n+1}, q_n, \ldots, q_1, q_0$. The points $p_1, \ldots, p_n$ represent $[n] \times \{0\}$ in the MIPA-instance, while the points $q_1, \ldots, q_n$ represent $[n] \times \{1\}$.

For every pair $i \leqslant j \in [n]$, we can associate a line $\ell_p(i,j)$ passing through middle(p_{i-1}, p_i) and middle(p_j, p_{j+1}). Notice that, by convexity, $\ell_p(i,j)$ separates the points $p_i, p_{i+1}, \ldots, p_{j-1}$, p_j (below it) from the points $p_1, \ldots, p_{i-1}$, $p_{j+1}, \ldots, p_n$ (above it). We similarly define $\ell_q(i,j)$ as the line passing through middle(q_{i-1}, q_i) and middle(q_j, q_{j+1}). We observe that $\ell_p(i,j)$ and $\ell_q(i,j)$ are parallel. For every interval $I = [i,j] \in \mathcal{I}$, we introduce in the maximum clique-instance the half-plane $h_p(I) := h_p(i,j)$ as the closed upper half-plane whose boundary is $\ell_p(i,j)$, and $h_q(I) := h_q(i,j)$ as the closed lower half-plane whose boundary is $\ell_q(i,j)$. We give these two objects weight 5 by superimposing 5 copies of them. All pairs of introduced half-planes intersect, except the pairs $\{h_p(i,j), h_q(i,j)\}$.

Finally for every edge $(i,0)(j,1)$ of the matching M (with $i,j \in [n]$), we add an axis-parallel rectangle $R(i,j)$ whose top-left corner is p_i and bottom-right corner is q_j. This finishes the construction of $(\mathcal{R}, \mathcal{H})$. When λ tends to 0, the rectangles are arbitrary close to squares of equal side-length. In other words, for any $\varepsilon > 0$, the axis-parallel rectangles can be made ϵ-squares. The half-planes can be turned into unit disks, making the side-length of the rectangles very small compared to 1. We denote by $(\mathcal{R}, \mathcal{D})$ the corresponding sets of axis-parallel rectangles and unit disks, and by G their intersection graph.

Let consider instances of MIPA produced by the previous reduction from Positive NAE 3-SAT-3, on ν-variable μ-clause formulas that are either satisfiable or with at least $(1-\gamma)\mu$ non satisfiable clauses. We call *yes-instances* the former MIPA-instances, and *no-instances*, the latter. If $(M, \mathcal{I})$ is a yes-instance, we claim that G has a clique of size $5|\mathcal{I}| + 3\nu + 4\mu$. Indeed there is a placement p that preserves $3\nu + 4\mu$ edges of M. We start by taking in the clique all the half-planes (or corresponding unit disks) $h_p(I)$ whenever $p(I) = 0$, and $h_q(I)$ whenever $p(I) = 1$. Since these objects have weight 5 (actually 5 stacked copies), this amounts to $5|\mathcal{M}|$ vertices. The corresponding half-planes pairwise intersect since their boundaries have distinct slopes. Then we include to the clique the $3\nu + 4\mu$ rectangles

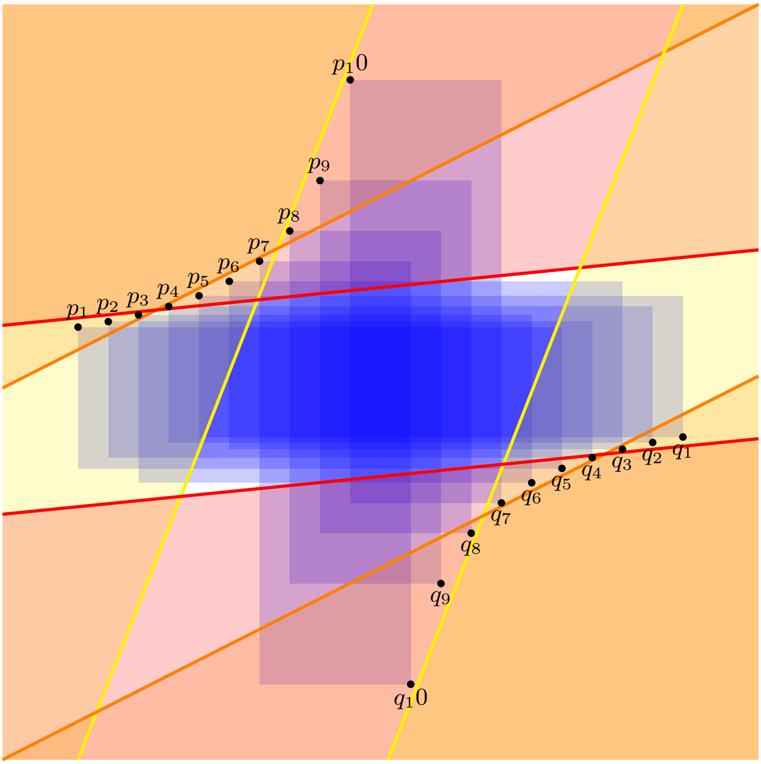

Figure 2.4: The output of the reduction on the instance of Figure 2.3.

$R(i,j) \in \mathcal{R}$ such that $(i,0)(j,1)$ is preserved by p. All the rectangles pairwise intersect since they all contain the origin O. Every pair of chosen half-plane $h_z(I)$ ($z \in \{p,q\}$) and rectangle $R(a,b)$ intersects, otherwise the placement of I would cover $(a,0)(b,1)$. Thus we exhibited a clique of size $5|\mathcal{I}| + 3\nu + 4\mu$ in G.

We now assume that $(M, \mathcal{I})$ is a no-instance, and we claim that G has no clique larger than $5|\mathcal{I}|+3\nu+4\mu-(1-\gamma)\mu$. Let us see how to build a clique in G. One can take at most one object between $h_p(I)$

and $h_q(I)$ (since they do not intersect). There is a maximum clique that takes at least one of $h_p(I)$ and $h_q(I)$ since $h_p(I)$ has weight 5 and intersects every object but $h_q(I)$ plus at most 5 rectangles (recall that the intervals of $\mathcal{I}$ have size at most 5). Thus we assume that our maximum clique takes exactly one object between $h_p(I)$ and $h_q(I)$, for every $I \in \mathcal{I}$. We consider the placement p defined as $p(I) = 0$ if $h_p(I)$ is in the clique, and $p(I) = 1$ if $h_q(I)$ is in the clique. Now the rectangles $R(i, j)$ that are adjacent to the chosen half-planes of $\mathcal{H}$ (or unit disks of $\mathcal{D}$) correspond to the edges $(i, 0)(j, 1)$ of M which are preserved. By Lemma 2.14, there are at most $3\nu + 4\mu - (1 - \gamma)\mu$ such rectangles.

Since $|V(G)| = |\mathcal{H}| + |\mathcal{R}| = 10|\mathcal{I}| + |M| = O(\nu + \mu) = O(\nu)$ and $\frac{5|I|+3\nu+4\mu-(1-\gamma)\mu}{|I|+3\nu+4\mu} \leqslant 1 - \frac{(1-\gamma)\mu}{140\mu+9\mu+4\mu} = 1 - \frac{1-\gamma}{153} = c$, by Theorem 2.12, maximum clique cannot be c-approximated in time $2^{|V(G)|^{1-\delta}}$, under the ETH. Besides, Corollary 2.13 implies that this problem cannot be 7633010347/7633010348-approximated in polynomial-time, unless P=NP. In particular, it is NP-hard and even APX-hard. □

Of course the ratios that are shown not achievable, even in time subexponential, under the ETH, are very close to 1. The current best exact exponential algorithm solving 3-SAT has running time 1.308^n [42], building upon the PPSZ algorithm [67]. Assuming getting this down to 1.14^n is impossible, which implies $s_3 > 0.2$, the inapproximability bound in subexponential-time of respectively γ' for Max Interval Permutation Avoidance and c for maximum clique in intersection graphs of unit disks (or half-planes) and axis-parallel rectangles are roughly $1 - 6 \cdot 10^{-26}$ and $1 - 5 \cdot 10^{-27}$, respectively.

Let us briefly discuss the issue the *co-2-subdivision* approach encounters for maximum clique in intersection graphs of half-planes and axis parallel-rectangles. Axis-parallel rectangles cannot represent a large antimatching (they already cannot represent $\overline{3K_2}$). Hence, as in our construction, the large antimatching has to be, for the most part, realised by half-planes. Now in the MIPA approach, the axis-parallel rectangles can avoid *two* arbitrary half-planes with the freedom of their top-left and bottom-right corners. In the *co-2-subdivision* approach, they would have to avoid *at least three* arbitrary half-planes, and do not have enough degrees of freedom for that.

2.4 2-pancakes and unit disks

In this section we prove the following theorem:

Theorem 2.16. *There exists a polynomial time algorithm for computing a maximum clique in Π^2, even without a representation.*

We first give a proof when a representation is given. The idea of the algorithm is similar to the one of Clark, Colbourn and Johnson [19]. We prove that if u and v are the most distant vertices in a maximum clique, then $\mathcal{N}(u) \cap \mathcal{N}(v)$ is cobipartite. In a second part, we give a robust algorithm, meaning that it does not require a representation, using the tools introduced by Raghavan and Spinrad [70].

2.4.1 Additional definitions and notations

For any $x_1 \leq x_2$, we denote by $P^2(x_1, x_2)$ the 2-pancake that is the Minkowski sum of the unit disk centred at O and the line segment with endpoints x_1 and x_2. Therefore we have $P^2(x_1, x_2) = \bigcup_{x_1 \leq x' \leq x_2} \mathcal{D}((x', 0), 1)$. Behind the definition of the d-pancakes is the idea that they should be as similar as possible to unit d-balls. In particular 2-pancakes should behave as much as possible like unit disks. This is perfectly illustrated when the intersection of a 2-pancake and a unit disk is a lens, as the intersection of two unit disks would be.

Definition 2.17. Let $\{P_j^2\}_{1 \leq j \leq n}$ be a set of 2-pancakes. For any unit disk $\mathcal{D}$, we denote by $L(\mathcal{D}, \{P_j^2\})$, or simply by $L(\mathcal{D})$ when there is no risk of confusion, the set of 2-pancakes in $\{P_j^2\}$ whose intersection with $\mathcal{D}$ is a lens.

Let $\mathcal{D}$ denote $\mathcal{D}(c,1)$ for some point c. Observe that if a 2-pancake $P^2(x_1,x_2)$ for some $x_1 \leq x_2$ is in $L(\mathcal{D})$, then the intersection between $\mathcal{D}$ and $P^2(x_1,x_2)$ is equal to $\mathcal{D} \cap \mathcal{D}((x_1,0),1)$ or $\mathcal{D} \cap \mathcal{D}((x_2,0),1)$. We make an abuse of notation and denote by $d(\mathcal{D},P^2(x_1,x_2))$ the smallest distance between c and a point in the line segment $[x_1,x_2]$. Observe that if the intersection between $\mathcal{D}$ and $P^2(x_1,x_2)$ is equal to $\mathcal{D}\cap\mathcal{D}((x_1,0),1)$, then $d(\mathcal{D},P^2(x_1,x_2)) = d(c,(x_1,0))$, and otherwise $d(\mathcal{D},P^2(x_1,x_2)) = d(c,(x_2,0))$. The following observation gives a characterisation of when the intersection between a unit disk and a 2-pancake is a lens.

Observation 2.18. Let $\mathcal{D}((c_x,c_y),1)$ be a unit disk intersecting with a 2-pancake $P^2(x_1,x_2)$. Then their intersection is a lens if and only if ($c_x \leq x_1$ or $c_x \geq x_2$) and the interior of $\mathcal{D}((c_x,c_y),1)$ does not contain any point in $\{(x_1,\pm 1),(x_2,\pm 1)\}$.

The observation follows immediately from the fact that the intersection is a lens if and only if $\mathcal{D}((c_x,c_y),1)$ does not contain a point in the open line segment between the points $(x_1,-1)$ and $(x_2,-1)$, nor in the open line segment between the points $(x_1,1)$ and $(x_2,1)$.

2.4.2 Maximum clique with a representation

In their proof, Clark, Colbourn and Johnson use the following fact: if c and c' are two points at distance ρ, then the diameter of the half-lenses induced by $\mathcal{D}(c,\rho)$ and $\mathcal{D}(c',\rho)$ is equal to ρ. We prove here a similar result.

Lemma 2.19. *Let c and c' be two points at distance ρ, and let be $\rho' \geq \rho$. Then the diameter of the half-lenses induced by $\mathcal{D}(c,\rho)$ and $\mathcal{D}(c',\rho')$ is at most ρ'.*

Proof. First note that if $\rho' > 2\rho$ then the half-lenses are half-disks of $\mathcal{D}(c,\rho)$. The diameter of these half-disks is equal to 2ρ, which is smaller than ρ'. Let us now assume that we have $\rho' \leq 2\rho$. The boundary of the lens induced by $\mathcal{D}(c,\rho)$ and $\mathcal{D}(c',\rho')$ consists of two arcs. The line (c,c') intersects exactly once with each arc. One of these two intersections is c', we denote by c'' the other. Let us consider the disk $\mathcal{D}(c'',\rho')$. Note that it contains the disk $\mathcal{D}(c,\rho)$. Therefore the lens induced by $\mathcal{D}(c,\rho)$ and $\mathcal{D}(c',\rho')$ is contained in the lens induced by $\mathcal{D}(c'',\rho')$ and $\mathcal{D}(c',\rho')$, whose half-lenses have diameter ρ'. The claim follows from the fact that the half-lenses of the first lens are contained in the ones of the second lens. □

Before stating the next lemma, we introduce the following definition:

Definition 2.20. Let $\{S_i\}_{1\leq i\leq n}$ and $\{S'_j\}_{1\leq j\leq n'}$ be two families of sets in $\mathbb{R}^2$. We say that $\{S_i\}$ and $\{S'_j\}$ *fully intersect* if for all $1 \leq i \leq n$ and $1 \leq j \leq n'$ the intersection between S_i and S'_j is not empty.

Lemma 2.21. *Let $\mathcal{D} := \mathcal{D}(c,1)$ be a unit disk and let the 2-pancake $P^2 := P^2(x_1,x_2)$ be in $L(\mathcal{D})$. Let $\{\mathcal{D}_i\}$ be a set of unit disks that fully intersect with $\{\mathcal{D},P^2\}$, such that for any unit disk $\mathcal{D}_i$ we have $d(\mathcal{D},\mathcal{D}_i) \leq d(\mathcal{D},P^2)$. Moreover if P^2 is in $L(\mathcal{D}_i)$ we require $d(\mathcal{D}_i,P^2) \leq d(\mathcal{D},P^2)$. Also let $\{P^2_j\}$ be a set of 2-pancakes that fully intersect with $\{\mathcal{D},P^2\}$, such that for any P^2_j in $\{P^2_j\} \cap L(D)$, we have $d(\mathcal{D},P^2_j) \leq d(\mathcal{D},P^2)$. Then $G(\{\mathcal{D}_i\} \cup \{P^2_j\})$ is cobipartite.*

Proof. The proof is illustrated in Figure 2.5. Without loss of generality, let us assume that the intersection between $\mathcal{D}$ and P^2 is equal to $\mathcal{D} \cap \mathcal{D}((x_1,0),1)$. Remember that by definition we have $x_1 \leq x_2$. Let $P^2(x'_1,x'_2)$ be a 2-pancake in $\{P^2_j\}$. As it is intersecting with P^2, we have $x'_2 \geq x_1 - 2$. Assume by contradiction that we have $x'_1 > x_1$. Then with Observation 2.18, we have that $P^2(x'_1,x'_2)$ is in $L(\mathcal{D})$ and

$d(\mathcal{D}, P^2(x_1', x_2')) > d(\mathcal{D}, P^2)$, which is impossible. Therefore we have $x_1' \leq x_1$, and so $P^2(x_1', x_2')$ must contain $\mathcal{D}((x', 0), 1)$ for some x' satisfying $x_1 - 2 \leq x' \leq x_1$. As the line segment $[(x_1 - 2, 0), (x_1, 0)]$ has length 2, the 2-pancakes in $\{P_j^2\}$ pairwise intersect.

We denote by ρ the distance $d(\mathcal{D}, P^2)$. Let $\mathcal{D}(c_i, 1)$ be a unit disk in $\{\mathcal{D}_i\}$. By assumption, c_i is in $\mathcal{D}(c, \rho) \cap \mathcal{D}((x_1, 0), 2)$. We then denote by R the lens that is induced by $\mathcal{D}(c, \rho)$ and $\mathcal{D}((x_1, 0), 2)$. We cut the lens into two parts with the line $(c, (x_1, 0))$, and denote by R_1 the half-lens that is not below this line, and by R_2 the half-lens that is not above it. With Lemma 2.19, we obtain that the diameter of R_1 and R_2 is at most 2. Let us assume without loss of generality that c is not below Ox. We denote by X_1 the set of unit disks in $\{\mathcal{D}_i\}$ whose centre is in R_1. We denote by X_2 the union of $\{P_j^2\}$ and of the set of unit disks in $\{\mathcal{D}_i\}$ whose centre is in R_2. Since the diameter of R_1 is 2, any pair of unit disks in X_1 intersect, therefore $G(X_1)$ is a complete graph. To show that $G(X_2)$ is a complete graph too, it remains to show that any unit disk $\mathcal{D}(c_i, 1)$ in X_2 and any 2-pancake $P^2(x_1', x_2')$ in $\{P_j^2\}$ intersect. We denote by P_+^2 the following convex shape: $\cup_{x_1' \leq x \leq x_2'} \mathcal{D}((x, 0), 2)$. Note that the fact that $\mathcal{D}(c_i, 1)$ and $P^2(x_1', x_2')$ intersect is equivalent to having c_i in P_+^2. Let us consider the horizontal line going through c, and let us denote by c' the left intersection with the circle centred at $(x_1, 0)$ with radius 2. We also denote by r_2 the extremity of R that is in R_2.

Let us assume by contradiction that c_i is above the line segment $[c, c']$. As by assumption c_i is in R_2, it implies that the x-coordinate of c_i is smaller than the one of c. Therefore P^2 is in $L(\mathcal{D}_i)$ and $d(\mathcal{D}_i, P^2) > d(\mathcal{D}, P^2)$, which is impossible by assumption. Let us denote by $R_{2,-}$ the subset of R_2 that is not above the line segment $[c, c']$. To prove that $\mathcal{D}(c_i, 1)$ and $P^2(x_1', x_2')$ intersect, it suffices to show that P_+^2 contains $R_{2,-}$. As shown above, $P^2(x_1', x_2')$ contains $\mathcal{D}((x', 0), 1)$ for some x' satisfying $x_1 - 2 \leq x' \leq x_1$. This implies that P_+^2 contains $\mathcal{D}((x_1 - 2, 0), 2) \cap \mathcal{D}((x_1, 0), 2)$, and in particular

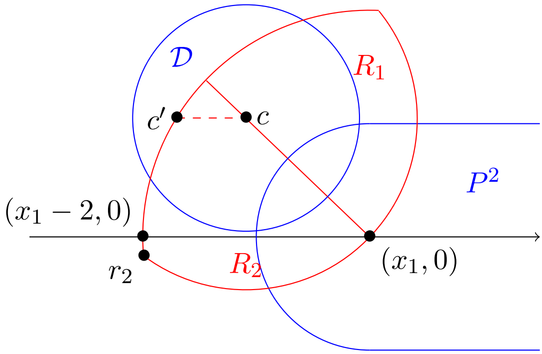

Figure 2.5: Illustration of the proof of Lemma 2.21

contains $(x_1, 0)$. Moreover as c is not below Ox, r_2 is also in $\mathcal{D}((x_1 - 2, 0), 2) \cap \mathcal{D}((x_1, 0), 2)$. As P^2 intersects $\mathcal{D}$, P^2_+ contains c. Let us assume by contradiction that P^2_+ does not contain c'. Then x'_2 must be smaller than the x-coordinate of c', because otherwise the distance $d((x'_2, 0), c')$ would be at most $d((x_1, 0), c')$, which is equal to 2. But then if P^2_+ does not contain c', then it does not contain c either, which is a contradiction. We have proved that P^2_+ contains the points $(x_1, 0)$, c, c' and r_2. By convexity, and using the fact that two circles intersect at most twice, we obtain that $R_{2,-}$ is contained in P^2_+. This shows that any two elements in X_2 intersect, which implies that $G(X_2)$ is a complete graph. Finally, as $X_1 \cup X_2 = \{\mathcal{D}_i\} \cup \{P^2_j\}$, we obtain that $G(\{\mathcal{D}_i\} \cup \{P^2_j\})$ can be partitioned into two cliques, i.e. it is cobipartite. □

Lemma 2.22. *Let $\mathcal{D} := \mathcal{D}((c_x, c_y), 1)$ and $\mathcal{D}' := \mathcal{D}((c'_x, c'_y), 1)$ be two unit disks such that $c_x \leq c'_x$. Let $P^2_1 := P^2(x_1, x_2)$ be a 2-pancake intersecting with $\mathcal{D}$ and $\mathcal{D}'$, such that $x_1 \geq c_x$ and P^2_1 is not in $L(\mathcal{D})$. If $P^2_2 := P^2(x'_1, x'_2)$ is a 2-pancake intersecting with $\mathcal{D}$ and $\mathcal{D}'$, but not intersecting with P^2_1, then P^2_2 is in $L(\mathcal{D}) \cap L(\mathcal{D}')$.*

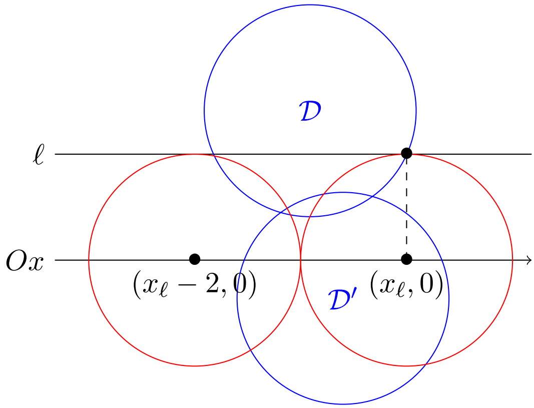

Figure 2.6: Illustration of the proof of Lemma 2.22

Proof. The proof is illustrated in Figure 2.6. First let us prove that P_2^2 cannot be on the right side of P_1^2, i.e. we have $x_1' \leq x_1$. Let us assume by contradiction $x_1' > x_1$. As P_1^2 and P_2^2 are not intersecting, we have $x_1' > x_1 + 2$. Hence, since we assume $x_1 \geq c_x$, we obtain $d(c, (x_1', 0)) > 2$, which is impossible. Therefore we have $x_1' \leq x_1$, and even $x_1' < x_1 - 2$ since P_1^2 and P_2^2 are not intersecting.

Without loss of generality, let us assume $c_y \geq 0$. Let us consider the horizontal line ℓ with height 1. By assumption it intersects with the circle centred at (c_x, c_y) with unit radius. There are at most two intersections, and we denote by x_ℓ the x-coordinate of the one to the right. As P_1^2 is not in $L(\mathcal{D})$, we have $x_1 \leq x_\ell$. Then, since $x_1' < x_1 - 2$ and the fact that $\mathcal{D}$ has diameter 2, we know that the points $(x_1', 1)$ and $(x_1', -1)$ are not in $\mathcal{D}$, which implies that P_2^2 is in $L(\mathcal{D})$. Likewise as we have $c_x \leq c_x'$, the points $(x_1', 1)$ and $(x_1', -1)$ are not in $\mathcal{D}'$, and so P_2^2 is in $L(\mathcal{D}) \cap L(\mathcal{D}')$. □

Lemma 2.23. *Let $\mathcal{D} := \mathcal{D}(c, 1)$ and $\mathcal{D}' := \mathcal{D}(c', 1)$ be two intersecting unit disks. Let $\{\mathcal{D}_i\}$ be a set of unit disks that fully intersect with*

$\{\mathcal{D}, \mathcal{D}'\}$, such that for each unit disk $\mathcal{D}_i$ we have $d(\mathcal{D}, \mathcal{D}_i) \leq d(\mathcal{D}, \mathcal{D}')$ and $d(\mathcal{D}', \mathcal{D}_i) \leq d(\mathcal{D}, \mathcal{D}')$. Also let $\{P_j^2\}$ be a set of 2-pancakes that fully intersect with $\{\mathcal{D}, \mathcal{D}'\}$, such that for any P_j^2 in $\{P_j^2\} \cap L(D)$, we have $d(\mathcal{D}, P_j^2) \leq d(\mathcal{D}, \mathcal{D}')$, and for any P_j^2 in $\{P_j^2\} \cap L(D')$, we have $d(\mathcal{D}', P_j^2) \leq d(\mathcal{D}, \mathcal{D}')$. Then $G(\{\mathcal{D}_i\} \cup \{P_j^2\})$ is cobipartite.

Proof. We denote by ρ the distance between c and c'. We denote by R the lens induced by $\mathcal{D}(c, \rho)$ and $\mathcal{D}(c', \rho)$. We cut R with the line segment $[c, c']$, which partitions R into two half-lenses that we denote by R_1 and R_2. By assumption, the centre of any unit disk in $\{\mathcal{D}_i\}$ must be in R. Since R_1 and R_2 have diameter ρ which is at most 2, any two unit disks having their centres in R_1 must intersect, and the same holds with R_2. Therefore $G(\{\mathcal{D}_i\})$ is cobipartite, which is the claim if $\{P_j^2\}$ is empty.

We now assume that $\{P_j^2\}$ is not empty. In order to show the claim, we do a case analysis according to whether the intersection between $\mathcal{D}(c, 2) \cap \mathcal{D}(c', 2)$ and Ox is empty or not. Let us assume that the latter holds, as shown in Figure 2.7. Let P^2 be in a 2-pancake in $\{P_j^2\}$. As P^2 intersects with $\mathcal{D}$, P^2 contains a unit disk that intersects with $\mathcal{D}$. Likewise, P^2 contains a unit disk that intersects with $\mathcal{D}'$. This implies that P^2 contains a point in $\mathcal{D}(c, 2) \cap Ox$ and a point in $\mathcal{D}(c', 2) \cap Ox$. By convexity of a 2-pancake, P^2 contains a point $(x', 0)$, where $(x', 0)$ is in $\mathcal{D}(c, 2) \cap \mathcal{D}(c', 2) \cap Ox$. We denote by R^+ the lens that is induced by $\mathcal{D}(c, 2)$ and $\mathcal{D}(c', 2)$ and cut it with the line (c, c'). We denote by R_1^+ (respectively R_2^+) the half-lens that contains R_1 (respectively R_2). Let us assume that $(x', 0)$ is in R_1^+. By assumption $\mathcal{D}((x', 0), 2)$ contains c and c'. Let us consider the third extremity of R_1, along with c and c', that we denote by r_1. By making a circle centred at r_1 grow, we observe that the farthest point from r_1 in R_1^+ can only be at one of the three extremities of R_1^+. However by Lemma 2.19 these distances are at most 2, which implies that the distance between $(x', 0)$ and r_1 is at most 2. Using the fact that two circles intersect at most twice, we obtain that R_1 is contained in $\mathcal{D}((x', 0), 2)$. Therefore P^2 intersect with all unit disks whose centre is in R_1, and with all 2-pancakes in $L(\mathcal{D}) \cap L(\mathcal{D}')$ that contain a disk whose centre is in R_1. Let $P^2(x_1, x_2)$ and $P^2(x'_1, x'_2)$ be two 2-pancakes in $\{P^2_j\}$ such that they contain each a unit disk whose centre is in R_1^+, but such that they do not contain a unit disk whose centre is in R_1. In particular, $P^2(x_1, x_2)$ and $P^2(x'_1, x'_2)$ are not in $L(\mathcal{D}) \cap L(\mathcal{D}')$. We claim that they intersect. Suppose by contradiction that they do not. Without loss of generality, let us assume that $P^2(x_1, x_2)$ is to the right of $P^2(x'_1, x'_2)$, and that $c_x \leq c'_x$, where c_x and c'_x denote the x-coordinate of c and c' respectively. Since $P^2(x_1, x_2)$ does not contain a disk in R_1, and since it is on the right side of $P^2(x'_1, x'_2)$, it implies that it does not contain a disk with centre in $\mathcal{D}(c, \rho)$. Therefore $P^2(x_1, x_2)$ cannot be in $L(\mathcal{D})$. Moreover the fact that it does not contain a disk with centre in $\mathcal{D}(c, \rho)$ implies $x_1 \geq c_x$. We finally apply Lemma 2.22 to obtain a contradiction. We denote X_1 the set of unit disks whose centre is in R_1 and 2-pancakes that contain a disk whose centre is in R_1^+. We know that two unit disks in X_1 intersect. Moreover we have shown that a 2-pancake and a unit disk in X_1 intersect. For a pair of two pancakes, if one of them contains a disk whose centre is in R_1 it is done for the same reasons. If none of them does, then we have shown above that they intersect. This shows that $G(X_1)$ is a complete graph. By defining X_2 as the set of the remaining disks and 2-pancakes, using the symmetry of the problem we obtain that $G(X_2)$ is also a complete graph.

Now let us assume that the intersection between $\mathcal{D}(c, 2) \cap \mathcal{D}(c', 2)$ and Ox is empty, as shown in Figure 2.8. As $\{P^2_j\}$ is not empty, the set $Ox \setminus (\mathcal{D}(c, 2) \cup \mathcal{D}(c', 2))$ consists of three connected component, one of them bounded. We denote by s the closed line segment consisting of the bounded connected component and its boundaries. Any 2-pancake P^2 in $\{P^2_j\}$ contains a unit disk whose centre is in $\mathcal{D}(c, 2) \cap Ox$, otherwise P^2 would not intersect with $\mathcal{D}$. Likewise P^2 contains a unit disk whose centre is in $\mathcal{D}(c', 2) \cap Ox$, and therefore P^2 contains

s. This implies that all 2-pancakes in $\{P^2_j\}$ pairwise intersect. Let us assume without loss of generality that R_1 is closer to Ox than R_2. Let us show that any 2-pancake P^2 in $\{P^2_j\}$ and any unit disk whose centre is in R_1 intersect. This is equivalent to show that for any point p in R_1, there exists a unit disk contained in P^2 with centre $q \in P^2 \cap Ox$, such that the Euclidean distance between p and q is at most 2. We denote by P^2_+ the Minkowski sum of the disk with radius 2 centred at O and the line segment s, i.e. $P^2_+ = \cup_{x' \in s} \mathcal{D}(x', 2)$. Note that P^2_+ is convex. We claim that P^2_+ contains R_1, which implies the desired property. Since s contains a point p_1 in $\mathcal{D}(c, 2)$, we know that P^2_+ contains c. Likewise, as s contains a point p_2 in $\mathcal{D}(c', 2)$, then P^2_+ contains c', and therefore by convexity the whole line segment $[c, c']$. Therefore P^2_+ contains the quadrilateral $cc'p_2p_1$. If this quadrilateral contains R_1 we are done. Otherwise, it may not contain a circular segment of the disk $\mathcal{D}(c', \rho)$ or a circular segment of the disk $\mathcal{D}(c, \rho)$. Let us assume that we have the worst case, meaning that both circular segments are not in $cc'p_2p_1$. Let us consider the circle $\mathcal{C}_1$ centred at p_1 with radius 2, and the circle $\mathcal{C}'$ centred at c' with radius ρ. The two circles intersect at c. Let us consider the point p'_1 that is at the intersection between $\mathcal{C}'$ and the line segment $[c, p_1]$. By definition, p'_1 is inside the disk induced by $\mathcal{C}_1$. As two circles intersect at most twice, we obtain that the arc cp'_1 centred at c' with radius ρ is contained in the disk induced by $\mathcal{C}_1$, and therefore also in P^2_+. By convexity, we know that the circular segment of the disk $\mathcal{D}(c', \rho)$ with chord $[c, p'_1]$ is in P^2_+. We can apply the same arguments for the other side to show that R_1 is in P^2_+. Hence by defining X_1 as the set of disks whose centre centre is in R_1, union the set of 2-pancakes, and X_2 as the set of disks whose centre is in R_2, we have that $G(X_1)$ and $G(X_2)$ are complete graphs. □

Note that Lemma 2.21 and Lemma 2.23 give a polynomial time algorithm for maximum clique in Π^2 when a representation is given. First compute a maximum clique that contains only 2-pancakes,

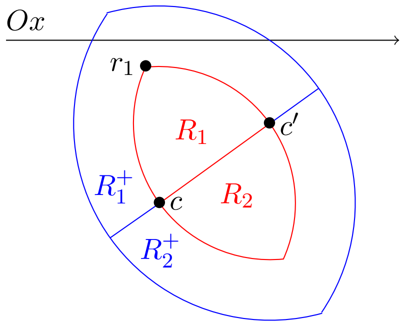

Figure 2.7: First case: $\mathcal{D}(c,2) \cap \mathcal{D}(c',2) \cap Ox \neq \emptyset$

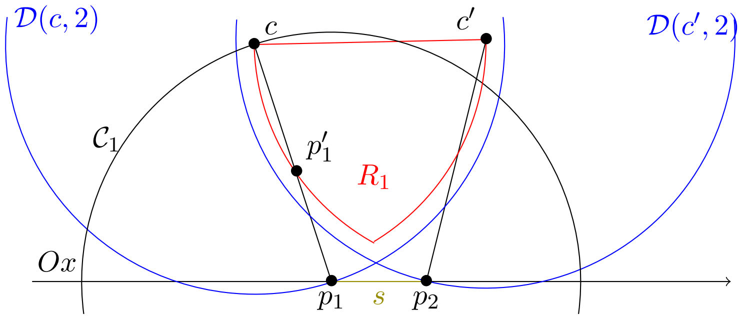

Figure 2.8: Second case: $\mathcal{D}(c,2) \cap \mathcal{D}(c',2) \cap Ox = \emptyset$

which can be done in polynomial time since the intersection graph of a set of 2-pancakes is an interval graph [39]. Then for each unit disk $\mathcal{D}$, compute a maximum clique which contains exactly one unit disk, $\mathcal{D}$, and an arbitrary number of 2-pancakes. Because finding out whether a unit disk and a 2-pancake intersect takes constant time, computing such a maximum clique can be done in polynomial time. Note that if a maximum clique contains at least two unit disks, then in quadratic time we can find in this maximum clique either a pair of unit disks or a unit disk and a 2-pancake whose intersection is a lens, such that the conditions of Lemma 2.21 or of Lemma 2.23 are satisfied. By applying the corresponding lemma, we know that we are computing a maximum clique in a cobipartite graph, which is the same as computing a maximum independent set in a bipartite graph. As this can be done in polynomial time [28], we can compute a maximum clique in Π^2 in polynomial time when the representation is given.

2.4.3 Maximum clique without a representation

To obtain an algorithm that does not require a representation, we use the notion of CNEEO as introduced by Raghavan and Spinrad [70]. The definition can be found in Section 2.1. To show that it is possible to compute a maximum clique in a graph G in Π^2, we show that such a graph admits a CNEEO.

Theorem 2.24. *If a graph G is in Π^2, then G admits a CNEEO.*

Theorem 2.24, Lemma 2.4 and Lemma 2.5 immediately imply Theorem 2.16. To prove Theorem 2.24, we use two more lemmas.

Lemma 2.25. *Let $\mathcal{D} = \mathcal{D}((c_x, c_y), 1)$ be a unit disk. Let $\{P^2_j\}$ be a set of 2-pancakes that all intersect with $\mathcal{D}$. Then $G(\{P^2_j\})$ is cobipartite.*

Proof. Let $P^2(x_1, x_2)$ be in $\{P^2_j\}$. By triangular inequality we have $x_1 \leq c_x + 2$ or $x_2 \geq c_x - 2$. It implies that $P^2(x_1, x_2)$ contains the line segment $[(x' - 1, 0), (x' + 1, 0)]$ for some x' satisfying $c_x - 2 \leq x' \leq c_x + 2$. We define X_1 as the set of 2-pancakes in $\{P^2_j\}$ that contain the line segment $[(x' - 1, 0), (x' + 1, 0)]$ for some x' satisfying $c_x - 2 \leq x' \leq c_x$, and X_2 as $\{P^2_j\} \setminus X_1$. We obtain that $G(X_1)$ and $G(X_2)$ are complete graphs. □

Lemma 2.26. *Let $P^2 = P^2(x_1, x_2)$ and $P'^2 = P^2(x'_1, x'_2)$ be two* *intersecting* 2*-pancakes. Let* $\{P_j^2\}$ *be a set of* 2*-pancakes that fully intersect with* $\{P^2, P'^2\}$*, such that for any* P_j^2 *in* $\{P_j^2\}$*,* P_j^2 *is not contained in* P^2 *nor in* P'^2*. Then* $G(\{P_j^2\})$ *is cobipartite.*

Proof. Let $P^2(x_1'', x_2'')$ be in $\{P_j^2\}$. Let us first assume that one of P^2, P'^2 is contained in the other. Without loss of generality, let us assume that P^2 is contained in P'^2, which is equivalent to having $x_1' \leq x_1 \leq x_2 \leq x_2'$. By assumption, as $P^2(x_1'', x_2'')$ is not contained in P^2, we have $x_1'' < x_1$ or $x_2 < x_2''$. As $P^2(x_1'', x_2'')$ intersects with P^2, it implies that $P^2(x_1'', x_2'')$ contains $(x_1 - 1, 0)$ or $(x_2 + 1, 0)$. We define X_1 as the set of 2-pancakes in $\{P_j^2\}$ that contains $(x_1 - 1, 0)$, and X_2 as $\{P_j^2\} \backslash X_1$. We obtain that $G(X_1)$ and $G(X_2)$ are complete graphs.

If none of P^2, P'^2 is contained in the other, we can assume without loss of generality that $x_1 \leq x_1' \leq x_2 \leq x_2'$. Therefore we have $x_1'' < x_1'$ or $x_2 < x_2''$, which implies that $P^2(x_1'', x_2'')$ contains $(x_1' - 1, 0)$ or $(x_2 + 1, 0)$. We conclude as above. □

Proof of Theorem 2.24. Let us consider any representation of G with unit disks and 2-pancakes. We divide the set of edges into three sets: E_1, E_2 and E_3. E_1 contains all the edges between a pair of unit disks, or between a unit disk $\mathcal{D}$ and a 2-pancake in $L(D)$. E_2 contains the edges between a unit disk and a 2-pancake that are not in E_1. E_3 contains the edges between a pair of 2-pancakes. For an edge $e = \{u, v\}$ in E_1, we call length of e the distance between u and v, be they unit disks or a unit disk $\mathcal{D}$ and a 2-pancake in $L(D)$. We order the edges in E_1 by non increasing length, which gives an ordering Λ_1. We take any ordering Λ_2 of the edges in E_2. For E_3, we take any ordering Λ_3 such that for any edge $e = \{u, v\}$, no edge after e in Λ_3 contains a 2-pancake contained in u or v. This can be obtained by considering the smallest 2-pancakes first. We finally define an ordering $\Lambda = \Lambda_1\Lambda_2\Lambda_3$ on E. Let us consider an edge e_k. If e_k is in E_1, Lemma 2.21 and Lemma 2.23 show that $N_{\Lambda,k}$ induces a cobipartite graph. If e_k is in E_2, we use Lemma 2.25, and if e_k is in E_3, we conclude with Lemma 2.26. This shows that Λ is a CNEEO. □

2.4.4 A motivation for Π^d

As we define it, Π^d is the class of intersection graphs of d-pancakes and unit d-balls. The properties that we desire are:

1. Π^d contains $(d-1)$-ball graphs and unit d-ball graphs,
2. Maximum clique can be computed as fast in Π^d as in $(d-1)$-ball graphs and unit d-ball graphs.

Let $\{\xi_i\}_{1 \leq i \leq d}$ be the canonical basis of $\mathbb{R}^d$. Let us consider another class $\tilde{\Pi}^d$, that might a priori satisfy those properties.

Definition 2.27. *We denote by $\tilde{\Pi}^d$ the class of intersection graphs of $(d-1)$-balls lying on the hyperspace induced by $\{\xi_1, \xi_2, \ldots, \xi_{d-1}\}$ and of unit d-balls.*

This class might look more natural since it makes use only of balls and not of pancakes. It contains by definition $(d-1)$-ball graphs and unit d-ball graphs. Moreover, as we want to be able to compute a maximum clique fast, we are looking for a "small" superclass. However, while we do not rule out the existence of a polynomial algorithm for computing a maximum clique in $\tilde{\Pi}^2$, we prove that Lemma 2.23 does not hold in $\tilde{\Pi}^2$.

The counterexample is illustrated in Figure 2.9. We have two intersecting unit disks $\mathcal{D}$ and $\mathcal{D}'$. Moreover each one of $\mathcal{D}_1$, $\mathcal{D}_2$ and the line segment $[x_1, x_2]$ intersects with both $\mathcal{D}$ and $\mathcal{D}'$. The distances $d(\mathcal{D}, \mathcal{D}_1)$, $d(\mathcal{D}', \mathcal{D}_1)$ are smaller than $d(\mathcal{D}, \mathcal{D}')$, and the same hold for $\mathcal{D}_2$. To define $L(\mathcal{D})$, a natural way would be to use the same characterisation as in Observation 2.18. Therefore the line segment $[x_1, x_2]$

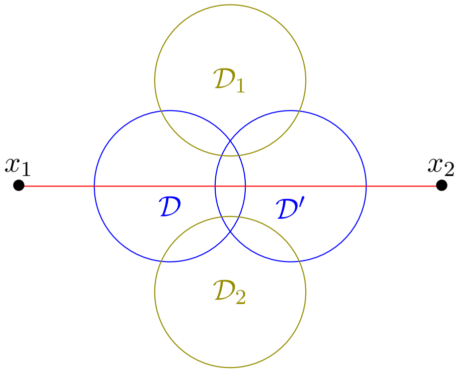

Figure 2.9: Lemma 2.23 does not hold in $\tilde{\Pi}^2$: $G(\{\mathcal{D}_1, \mathcal{D}_2, [x_1, x_2]\})$ is not cobipartite

is not in $L(\mathcal{D})$ nor in $L(\mathcal{D}')$. However, $G(\{\mathcal{D}_1, \mathcal{D}_2, [x_1, x_2]\})$ is an edgeless graph with three vertices, and therefore is not cobipartite.

2.4.5 Recognition of graphs in Π^2

We show that testing whether a graph can be obtained as the intersection graph of unit disks and 2-pancakes is hard, as claimed in the following theorem:

Theorem 2.28. *Testing whether a graph is in Π^2 is NP-hard, and even $\exists\mathbb{R}$-hard.*

Proof. We do a reduction from recognition of unit disk graphs, which is $\exists\mathbb{R}$-hard as shown by Kang and Müller [49]. Let $G = (V, E)$ be a graph with n vertices. We are going to construct $\binom{n}{2}$ graphs such that G is a unit disk graph if and only if at least one of these new graphs is in Π^2. Let S and S' be two stars with internal vertex s and s' respectively, having $14n + 8$ leaves each. Let W and W' be two paths with $2n$ vertices each with end vertices w_1, w_{2n} and w'_1, w'_{2n} respectively. Let u and v be two vertices in V. We define $G_{u,v}$ as the graph obtained by connecting s to s', w_1 to u, w'_1 to v, w_{2n} to s and w'_{2n} to s'. We claim that G is a unit disk graph if and only if $G_{u,v}$ is in Π^2 for some u and v in V. First let us assume that G is a unit disk graph. Let us consider the set P of the centres of the unit disks in any fixed representation of G. Consider two extreme points in P, meaning that removing any of them modifies the convex hull of the point set. Take the two unit disks $\mathcal{D}$ and $\mathcal{D}'$ corresponding to those two extreme points, and let us denote by u and v the corresponding vertices in G. Now take two sets $\{\mathcal{D}_i\}_{1\leq i\leq 2n}$ and $\{\mathcal{D}'_j\}_{1\leq j\leq 2n}$ of $2n$ unit disks such that $G(\{\mathcal{D}_i\})$ and $G(\{\mathcal{D}'_j\})$ are paths, and such that no two unit disks of the form $\mathcal{D}_i, \mathcal{D}'_j$ intersect. Moreover we require that $G(\{\mathcal{D}_i\}) \cap G = (\{u\}, \emptyset)$ and $G(\{\mathcal{D}_j\}) \cap G = (\{v\}, \emptyset)$, and that all unit disks centres are on the same side of the line (c_{2n}, c'_{2n}), which are the centres of $\mathcal{D}_{2n}$ and $\mathcal{D}'_{2n}$ respectively. This is possible because the most distant points in the unit disk representation of G have distance at most $4n$, and we have $2n$ unit disks in each path. Then we translate and rotate everything so that the y-coordinate of c_{2n} and c'_{2n} is equal to 2, and that all other centres are above the horizontal line with height 2. We take two intersecting 2-pancakes such that one also intersect with $\mathcal{D}_{2n}$ and the other with $\mathcal{D}'_{2n}$. We choose these 2-pancakes big enough so that for each of them we can add $14n + 8$ pairwise non intersecting unit disks, but intersecting with their respective 2-pancake. This shows that if G is a unit disk graph, then $G_{u,v}$ is in Π^2.

Let us now assume that $G_{u,v}$ is in Π^2, for some u, v in V. As a unit disk can intersect at most with 5 pairwise non intersecting unit disks, we have that in any Π^2 representation of $G_{u,v}$, s and s' must be represented by 2-pancakes, denoted by P and P' respectively. Let x be the length of the line segment obtained as the intersection of P and Ox. Note that all points of a unit disk intersecting a 2-pancake are within distance 3 of Ox. Therefore, the unit disks corresponding to leaves of s are contained in a rectangle with area $6(x + 4)$. Moreover, for each 2-pancake intersecting P, there is a unit disk contained in this 2-pancake that intersects P. We have $14n+8$ pairwise non-intersecting unit disks in a rectangle with area $6(x+4)$. As the area of a unit disk is bigger than 3, we have $6(x+4) \geq 3(14n+8)$, or equivalently $x \geq 7n$. Note that the same holds with P'. Let us show that in any Π^2 representation of $G_{u,v}$, all the vertices in V are represented by unit disks. Assume by contradiction that it is not the case. Without loss of generality, let us assume P is to the left of P', and that one vertex u_G in V is represented by a 2-pancake that is to the right of P'. Indeed this 2-pancake cannot be between P and P' because they are intersecting. Let us consider the last vertex in a path from s to u_G that is a disk. By construction, the distance between P and the unit disk corresponding to this vertex is at most $2(2n+n-1) = 6n-2$. This shows that this vertex is still far from the right end of P', and so the next vertex has to be represented by a unit disk because it is not intersecting P', which is a contradiction. We have shown that G is a unit disk graph if and only if there exist u, v in V such that $G_{u,v}$ is in Π^2, and the construction of these $\binom{n}{2}$ graphs takes linear time for each of them. □

CHAPTER 3

Maximum Clique in superclasses of disk graphs

3.1 Introduction

In Chapter 2, we studied maximum clique in superclasses of unit disk graphs. In Section 2.2, we extended the polynomial-time algorithms by Clark, Colbourn and Johnson [19] and Raghavan and Spinrad [70] to the case of translates of a fixed bounded convex set. In Section 2.4, we extended it to Π^2, the class of intersection graphs of unit disks and 2-pancakes. In this chapter, we study superclasses of disk graphs, and try to extend the EPTAS of Bonamy *et al.* [8] to new settings. We first give some definitions.

A cycle is said *induced* if it is chordless. An *odd cycle* (respectively *even cycle*) is a cycle on an odd (respectively even) number of vertices. One can observe that an odd cycle always contains an induced odd cycle. Two cycles are said *mutually induced* if they are chordless and there is no edge linking a vertex of one to a vertex of the other. The *induced odd cycle packing number* is the maximum number of disjoint odd cycles, that are pairwise mutually induced. We denote the induced odd cycle packing number of a graph G by $\text{iocp}(G)$.

Bonamy *et al.* showed:

Theorem 3.1 ([8]). *For any constants $d \in \mathbb{N}$, $0 < \beta \leqslant 1$, for every $0 < \varepsilon < 1$, there is a randomized $(1-\varepsilon)$-approximation algorithm running in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$, and a deterministic PTAS running in time $n^{\tilde{O}(1/\varepsilon^3)}$ for maximum clique in n-vertex graphs G satisfying the following conditions:*

- *there is a constant K such that $iocp(\overline{G}) \leq K$,*
- *the VC-dimension of the neighborhood hypergraph $\{N[v] \mid v \in V(G)\}$ is at most d, and*
- *G has a clique of size at least βn.*

The first item is enough to obtain a subexponential-algorithm [9] and boils down to proving a structural lemma on the representation of $K_{2,2}$ (see Lemma 3.6). Observe that in this statement we consider the complement of G, and not G itself. For intersection graphs, the first item is usually the hardest part to prove. For a disk graph or a unit ball graph G, Bonnet *et al.* and Bonamy *et al.* proved that we have $\text{iocp}(\overline{G}) \leq 1$ [9, 8]. Not all disk graphs satisfy the third item. However Bonnet *et al.* showed how to reduce the problem of maximum clique in disk graphs to the case where the third item is satisfied.

We show how to compute a maximum clique in homothets of a fixed centrally symmetric convex set using Theorem 3.1. Indeed we prove that for such a graph G, we have $\text{iocp}(\overline{G}) \leq 1$. We conjecture that the same holds with intersection graphs of convex-pseudo disks. We denote the class of convex pseudo-disk graphs by $\mathcal{G}$. The proof of Bonnet *et al.*, and ours for homothets of a fixed centrally symmetric convex set rely heavily on the fact that such objects have centres [9]. However, convex pseudo-disks do not, therefore adapting the proof in this new setting does not seem easy. While we are not able to extend this structural result to the class $\mathcal{G}$, we show a weaker property: The complement of a triangle and an odd cycle is a forbidden induced subgraph in $\mathcal{G}$. We write "complement of a triangle" to make the connection with iocp clear, but note that actually the complement of a triangle is an independent set of three vertices. Below we state this property more explicitly.

Theorem 3.2. *Let G be in $\mathcal{G}$. If there exists an independent set of size 3, denoted by H, in G, and if for any $u \in H$ and $v \in G \setminus H$, the edge $\{u, v\}$ is an edge of G, then $G \setminus H$ is cobipartite.*

Note that a cobipartite graph is not the complement of an odd cycle. Given the three pairwise non-intersecting convex pseudo-disks in H, we give a geometric characterisation of the two independent sets in the complement of $G \setminus H$. We conjecture that Theorem 3.2 is true even when H is the complement of any odd cycle, which implies:

Conjecture 3.3. *For any convex pseudo-disk graph G, we have $iocp(\overline{G}) \leq 1$.*

If Conjecture 3.3 holds, it is straightforward to obtain an EPTAS for maximum clique in convex pseudo-disks graphs, by using the method of Bonamy *et al.* [8].

Recall that Bonamy *et al.* asked for a geometric superclass of both disk graphs and unit ball graphs, in which maximum clique would admit an EPTAS. Let us consider the class Π^3, the class of intersection graphs of 3-pancakes and unit balls. We show that the following conjecture implies the existence of an EPTAS by using Theorem 3.1.

Conjecture 3.4. *There exists an integer K such that for any graph G in Π^3, we have $iocp(\overline{G}) \leq K$.*

3.2 Homothets of a fixed set

Here we observe that the EPTAS for maximum clique in disk graphs extends to the intersection graphs of homothets of a centrally symmetric convex set.

Theorem 3.5. *Maximum clique admits a subexponential-time algorithm and an EPTAS in intersection graphs of homothets of a fixed bounded centrally symmetric convex set S.*

We use the associated norm as defined in Section 2.2, and check the three conditions of Theorem 3.1.

Lemma 3.6. *In a representation of $K_{2,2}$ with homothets of S placing the four centers in convex position, the non-edges are between vertices corresponding to opposite corners of the quadrangle.*

Proof. Let S_1, S_2, S_3 and S_4 be the four homothets. We denote by c_i the center of S_i, and by λ_i its scaling factor. Let us assume by contradiction that they appear in this order on the convex hull, that S_1 and S_2 make one non-edge, and that S_3 and S_4 make the other. By assumption, we have $\|c_1 - c_2\| > \lambda_1 + \lambda_2$, and likewise $\|c_3 - c_4\| > \lambda_3 + \lambda_4$. Let us denote by c the intersection of the lines $\ell(c_1, c_3)$ and $\ell(c_2, c_4)$. We have $\|c_1 - c\| + \|c - c_2\| > \|c_1 - c_2\|$ by triangular inequality. Likewise it holds $\|c_3 - c\| + \|c - c_4\| > \|c_3 - c_4\|$. We therefore obtain $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 < \|c_1 - c\| + \|c - c_2\| + \|c_3 - c\| + \|c - c_4\| = \|c_1 - c_3\| + \|c_2 - c_4\| \leqslant \lambda_1 + \lambda_3 + \lambda_2 + \lambda_4$, which is a contradiction. $\square$

Lemma 3.6 implies by some parity arguments that the first condition of Theorem 3.1 holds (see Theorem 6 in [9]). It is well known that a

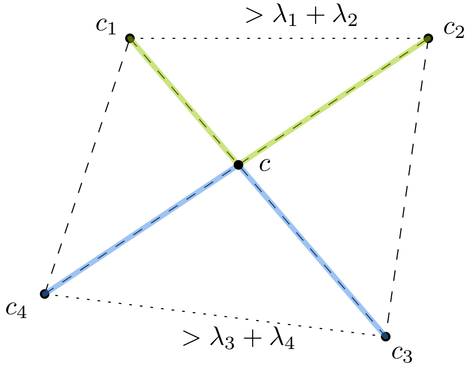

Figure 3.1: Illustration of the proof of Lemma 3.6. Non-edges are dotted and edges are dashed. By Assumption the top and bottom segment have a length of at least $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4$. By the triangle inequality the green plus the blue path are even longer.

family of homothets forms a pseudo-disk arrangement. Therefore the second property holds as shown by Aronov et al. [4]. Finally we enforce the third condition of Theorem 3.1, by using a chi-boundedness result of Kim et al. [51].

Lemma 3.7. *With a polynomial multiplicative factor in the running time, one can reduce to instances satisfying the third condition of Theorem 3.1 with* $\beta = 1/36$.

Proof. Kim et al. [51] show that in any representation of an intersection graph G of homothets of a convex set, a homothet S with a smallest scaling factor has degree at most $6\omega(G) - 7$, where $\omega(G)$ denotes the clique number of G. Their proof also implies that the independence number of its neighborhood is at most 6. By a degeneracy argument, the colouring number, denoted by $\chi(G)$ is at most $6\omega(G) - 6$. First we find in polynomial-time a vertex v such that the independence number of its neighborhood is at most 6. Let us denote by G_v the subgraph induced by its neighborhood, and n denotes its number of vertices. We denote by $\alpha(.)$ the independence number of a graph. As G_v has a representation with homothets of S, we have $\chi(G_v) \leqslant 6\omega(G_v)$. Therefore $\alpha(G_v)\omega(G_v) \geqslant \frac{1}{6}\alpha(G_v)\chi(G_v) \geqslant \frac{1}{6}n$. Thus by assumption we have $\omega(G_v) \geqslant \frac{1}{36}n$. Then we can compute a maximum clique that contains v, or remove v from the graph and iterate. The EPTAS of Bonamy et al. is called linearly many times. □

3.3 Convex pseudo-disks

In this section we are interested in computing a maximum clique in intersection graphs of convex pseudo-disks. As mentioned in the introduction, there exists an EPTAS for maximum clique in disk graphs [9, 8]. The main property used is that for any disk graph G, we have $\text{iocp}(\overline{G}) \leq 1$. The proof of this inequality relies on the fact that disks have centres. However in this section we are considering convex pseudo-disks, which do not have centres. Our proof that the complement of a triangle and an odd cycle is a forbidden induced subgraph in convex pseudo-disk graphs relies on line transversals and their geometric permutations on the three convex pseudo-disks that form a triangle in the complement, denoted by $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. As there are only three sets, the geometric permutation of a line transversal is given simply by stating which set is the second one intersected. We denote by $\{\mathcal{D}'_j\}_{1 \leq j \leq n}$ (or simply $\{\mathcal{D}'_j\}$) a set of convex pseudo-disks that fully intersect with $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$. Our aim is to show that $G(\{\mathcal{D}'_j\})$ is cobipartite.

Throughout this section, for the sake of readability, we refer to the convex pseudo-disks simply as “disks”. We always assume that no disk $\mathcal{D} \in \{\mathcal{D}'_j\}$ contains any of $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. Indeed, $\mathcal{D}$ would intersect any disk that intersects pairwise with $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. Thus, such a disk $\mathcal{D}$ could be arbitrarily added to any of the two cliques of the cobipartition.

Definition 3.8. A *line transversal* ℓ is a line that intersects each of the three disks $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. We call *disk in the middle* of a line transversal the disk it intersects in second position.

We are going to conduct a case analysis depending on the number of disks being the disk in the middle for some line transversal. If there exists no line transversal, we can prove a stronger statement.

Lemma 3.9. *If there is no line transversal through a family of convex sets F, then for any pair of convex sets $\{C_1, C_2\}$ that fully intersects with F, the sets C_1 and C_2 intersect.*

Proof. Let us prove the contrapositive. Assume that C_1 and C_2 do not intersect, therefore there exists a separating line. As all sets in F intersect C_1 and C_2, they also intersect the separating line, which is thus a line transversal of F. □

Using the notation of Theorem 3.2, Lemma 3.9 immediately implies that if there is no line transversal through the sets representing H, then $G \setminus H$ is a clique, which is an even stronger statement than required.

Definition 3.10. Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be two disjoint disks and let p, q be in the interior of $\mathcal{D}_1, \mathcal{D}_2$ respectively. We call *external tangents* of $\mathcal{D}_1$ and $\mathcal{D}_2$ the two tangents that do not cross the line segment $[p, q]$.

Definition 3.11. Let us consider a disk in $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$, say $\mathcal{D}_2$, such that it is the disk in the middle of a line transversal. We denote by τ and τ' the two external tangents of $\mathcal{D}_1$ and $\mathcal{D}_3$. We say that $\mathcal{D}_2$ is *contained* if it is contained in the bounded region S delimited by $\mathcal{D}_1$, τ, $\mathcal{D}_3$ and τ'. If $\mathcal{D}_2$ intersects exactly one of the external

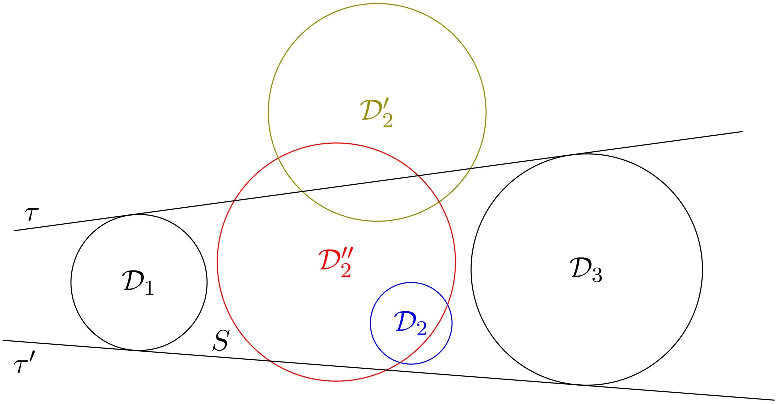

Figure 3.2: $\mathcal{D}_2$ is contained, $\mathcal{D}_2'$ is 1-intersecting and $\mathcal{D}_2''$ is 2-intersecting.

tangents, we say that $\mathcal{D}_2$ is 1-*intersecting*. If $\mathcal{D}_2$ intersects both external tangents, we say that $\mathcal{D}_2$ is 2-*intersecting*. The different cases are illustrated in Figure 3.2.

Lemma 3.12. *If $\mathcal{D}_2$ is 2-intersecting, then it is the disk in the middle of all line transversals.*

Proof. By definition, $\mathcal{D}_2$ is the disk in the middle of a line transversal. We denote by τ and τ' the external tangents. let p be a point in $\mathcal{D}_2 \cap \tau$ and p' be in $\mathcal{D}_2 \cap \tau'$. The line segment $[p, p']$ is contained in $\mathcal{D}_2$, and the line (p, p') separates $\mathcal{D}_1$ from $\mathcal{D}_3$. Let ℓ be a line transversal. Let p_1 be in $\ell \cap \mathcal{D}_1$ and p_3 be in $\ell \cap \mathcal{D}_3$. The line segment $[p_1, p_3]$ must cross $[p, p']$, which shows that the disk in the middle of ℓ is $\mathcal{D}_2$. ☐

Lemma 3.13. *If $\mathcal{D}_2$ is contained, then either $\mathcal{D}_1$ is not the disk in the middle of a line transversal, or $\mathcal{D}_1$ is 1-intersecting. The same holds with $\mathcal{D}_3$.*

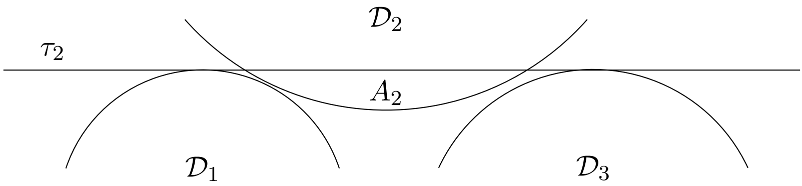

Figure 3.3: Illustration of Definition 3.14, with $\mathcal{D}_2$ being the 1-intersecting disk

Proof. We prove the statement for $\mathcal{D}_1$, the same arguments hold with $\mathcal{D}_3$. Observe that there is a line transversal having $\mathcal{D}_2$ as disk in the middle. Let us assume that there is a line transversal having $\mathcal{D}_1$ as disk in the middle, and let us show that $\mathcal{D}_1$ is 1-intersecting. As $\mathcal{D}_2$ is contained, no point in $\mathcal{D}_2$ lies on the boundary of the convex hull of $\mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3$. This immediately implies that some points in $\mathcal{D}_1$ lies on the boundary of the convex hull of $\mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3$. Therefore, $\mathcal{D}_1$ is not contained. Moreover, $\mathcal{D}_1$ is not 2-intersecting, for otherwise $\mathcal{D}_2$ would not be the disk in the middle of a line transversal, as stated in Lemma 3.12. We have shown that $\mathcal{D}_1$ is 1-intersecting. □

The following definition is illustrated in Figures 3.3, 3.4 and 3.5.

Definition 3.14. Let us consider a disk $\mathcal{D}_i$ in $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$ that is 1-intersecting, say $\mathcal{D}_i = \mathcal{D}_2$. We denote by τ_2 the external tangent of $\mathcal{D}_1$ and $\mathcal{D}_3$ that $\mathcal{D}_2$ intersects. We denote by A_2 the part of $\mathcal{D}_2$ that is on the same side of τ_2 as $\mathcal{D}_1$ and $\mathcal{D}_3$. Let $\mathcal{D}'$ be a disk intersecting pairwise with $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. We say that $\mathcal{D}'$ is *outside-containing* $\mathcal{D}_2$ if $\mathcal{D}_2 \setminus A_2$ is a subset of $\mathcal{D}'$. We denote by χ_1 and χ_2 the points where the boundaries of $\mathcal{D}'$ and $\mathcal{D}_2$ intersect. Note that they are both in A_2. We denote by $\mathcal{H}$ the closed halfplane with bounding line (χ_1, χ_2) that contains $\mathcal{D}_2 \setminus A_2$. Let $\mathcal{H}'$ be the closed halfplane with bounding line τ_2 that contains A_2. Note that $(\mathcal{H} \cap \mathcal{H}') \setminus A_2$ is the union of one or two connected sets. We have $\mathcal{D}' \cap \mathcal{D}_1 \subset (\mathcal{H} \cap \mathcal{H}') \setminus A_2$

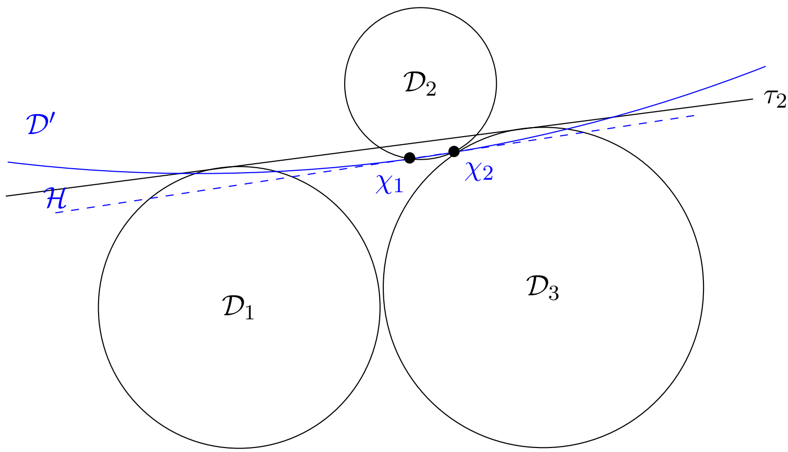

Figure 3.4: Illustration of Definition 3.14, $\mathcal{D}'$ is centred with respect to $\mathcal{D}_2$.

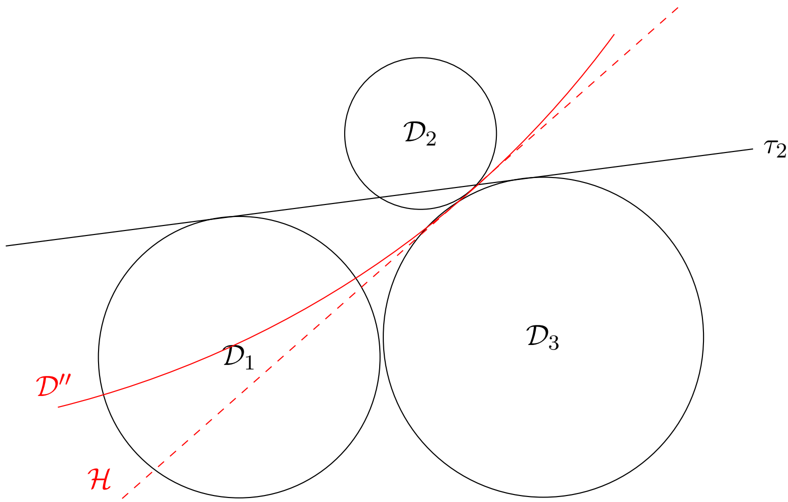

Figure 3.5: Illustration of Definition 3.14, $\mathcal{D}''$ is outside-containing $\mathcal{D}_2$, but not centred with respect to $\mathcal{D}_2$.

and $\mathcal{D}' \cap \mathcal{D}_3 \subset (\mathcal{H} \cap \mathcal{H}') \setminus A_2$. If $\mathcal{D}' \cap \mathcal{D}_1$ and $\mathcal{D}' \cap \mathcal{D}_3$ are not in the same connected set, we say that $\mathcal{D}'$ is *centred with respect to* $\mathcal{D}_2$.

Let us consider a disk in $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$, say $\mathcal{D}_2$, and let us assume it is the disk in the middle of a line transversal. Let us denote by τ and τ' the two external tangents of $\mathcal{D}_1$ and $\mathcal{D}_3$. Let $\mathcal{D}$ be a disk in $\{\mathcal{D}'_j\}$, such that there exists a pair of points $p_1 \in \mathcal{D} \cap \mathcal{D}_1$ and $p_3 \in \mathcal{D} \cap \mathcal{D}_3$, such that the segment $[p_1, p_3]$ intersects $\mathcal{D}_2$, potentially at a single point. Without loss of generality, we can even assume that $[p_1, p_3] \cap \mathcal{D}_1 = \{p_1\}$ and $[p_1, p_3] \cap \mathcal{D}_3 = \{p_3\}$. The segment $[p_1, p_3]$ splits $\mathcal{D}_2$ into two closed parts. One of them can potentially be a single point p if $[p_1, p_3] \cap \mathcal{D}_2 = \{p\}$. Observe that $\mathcal{D}$ contains exactly one of those two parts: at least one because of the pseudo-disk property, and at most one because $\mathcal{D}$ does not contain $\mathcal{D}_2$. Let us denote by $A_2 \subset \mathcal{D}_2$ the part that is contained in $\mathcal{D}$. We say that the side of $[p_1, p_3]$ where A_2 lies is the *positive side* of $[p_1, p_3]$. By definition, p_1 and p_3 lie between τ and τ'. By moving p_1 and p_3 toward the positive side of $[p_1, p_3]$ and along the boundary of $\mathcal{D}_1$ and $\mathcal{D}_3$ respectively, only two things can happen by construction: Either both of them reach τ, or both of them reach τ'.

Definition 3.15. We denote by $X_{\mathcal{D}_2}$ (or simply by X when there is no risk of confusion) the set of disks in $\{\mathcal{D}'_j\}$ for which the external tangent reached is τ. Likewise, we denote by $X'_{\mathcal{D}_2}$ (or simply by X') the set of disks in $\{\mathcal{D}'_j\}$ for which the external tangent reached is τ'.

Let us now assume the existence of a disk $\mathcal{D}' \in \{\mathcal{D}'_j\}$, which is not in $X \cup X'$. Let us consider $p_1 \in \mathcal{D}' \cap \mathcal{D}_1$ and $p_3 \in \mathcal{D}' \cap \mathcal{D}_3$. By assumption, the segment $[p_1, p_3]$ does not intersect $\mathcal{D}_2$. Observe that this implies that $\mathcal{D}_2$ is not 2-intersecting. It is possible to continuously move p_1 and p_3 in $\mathcal{D}_1$ and $\mathcal{D}_3$ respectively, such that they both reach either τ or τ', and while maintaining the property that $[p_1, p_3] \cap \mathcal{D}_2 = \emptyset$. Observe that the choice of $p_1 \in \mathcal{D} \cap \mathcal{D}_1$ and $p_3 \in \mathcal{D} \cap \mathcal{D}_3$ has no impact on whether they can both reach τ, or both reach τ'. Otherwise, it would be possible to move them from τ to τ' without having $[p_1, p_3]$ intersecting $\mathcal{D}_2$, which would imply that $\mathcal{D}_2$ is not the disk in the middle of any line transversal. If $\mathcal{D}_2$ is 1-intersecting, then exactly one of τ, τ' will always be reached, for any disk in $\{\mathcal{D}'_j\}$.

Definition 3.16. We denote by $Y_{\mathcal{D}_2}$ (or simply by Y when there is no risk of confusion) the set of disks in $\{\mathcal{D}'_j\} \setminus (X_{\mathcal{D}_2} \cup X'_{\mathcal{D}_2})$ for which the external tangent reached is τ, and which are not centred with respect to $\mathcal{D}_1$ or $\mathcal{D}_3$. Likewise, we denote by $Y'_{\mathcal{D}_2}$ (or simply by Y') the set of disks for which the external tangent reached is τ', and which are not centred with respect to $\mathcal{D}_1$ or $\mathcal{D}_3$. We denote by $Z_{\mathcal{D}_2}$ (or simply by Z), the set of disks for which the external tangent reached is τ', and which are centred with respect to $\mathcal{D}_1$ or $\mathcal{D}_3$. Finally, we denote by $Z'_{\mathcal{D}_2}$ (or simply by Z'), the set of disks for which the external tangent reached is τ, and which are centred with respect to $\mathcal{D}_1$ or $\mathcal{D}_3$.

We want to emphasise the fact that indeed in the definition of Z, p_1 and p_3 can reach τ' and not τ. This choice of notation comes from the fact that, assuming that $\mathcal{D}_1$ and $\mathcal{D}_3$ are the disks in the middle of no line transversal, or that they are 1-intersecting, all pairs of disks in $X \cup Y \cup Z$ intersect, and the same holds with all pairs of disks in $X' \cup Y' \cup Z'$, as we show later. Observe that if a disk is centred with respect to $\mathcal{D}_1$, then $\mathcal{D}_1$ is 1-intersecting. Thus if both $\mathcal{D}_1$ and $\mathcal{D}_3$ are not 1-intersecting, the sets Z and Z' are empty.

Lemma 3.17. *If $\mathcal{D}$ and $\mathcal{D}'$ are in $X_{\mathcal{D}_2}$, then they intersect.*

Proof. Let p_1, p_3, p'_1, p'_3 be points coming from the definition of $\mathcal{D}$ and $\mathcal{D}'$ being in X. Recall that p_1 and p'_1 lie on the boundary of $\mathcal{D}_1$. Similarly, p_3 and p'_3 lie on the boundary of $\mathcal{D}_3$. If $[p_1, p_3]$ and $[p'_1, p'_3]$ intersect then we are done. Otherwise, we can assume without loss of generality that p_1 is closer to τ than p'_1 is, and that p_3 is closer to

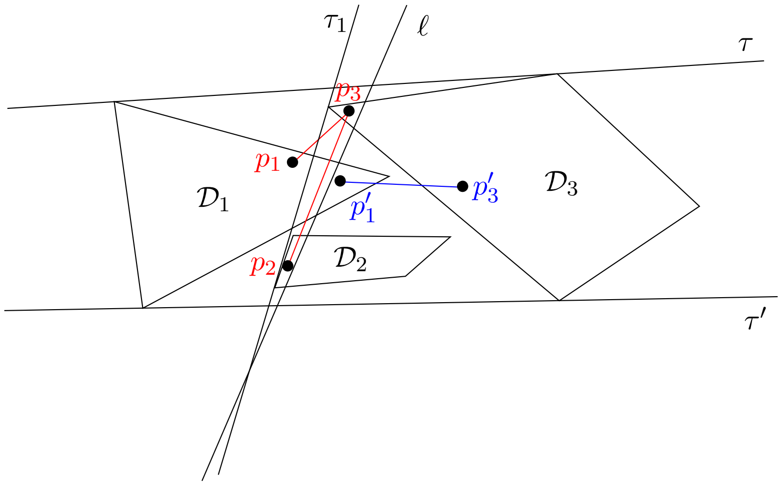

Figure 3.6: Illustration of Lemma 3.18. The segment $[p_2, p_3]$ splits $\mathcal{D}_1$, thus $\mathcal{D}$ contains p'_1 or is centred with respect to $\mathcal{D}_1$.

τ than p'_3 is (when considering them as points on the boundaries of $\mathcal{D}_1$ and $\mathcal{D}_3$ respectively). Let A_2, respectively A'_2, be the part of $\mathcal{D}_2$ contained in $\mathcal{D}$, respectively in $\mathcal{D}'$. By assumption, A_2 is contained in A'_2, which implies that the two disks intersect. $\square$

Lemma 3.18. *Let us assume that $\mathcal{D}_1$ is the disk in the middle of no line transversal or is* 1*-intersecting, and that the same holds with $\mathcal{D}_3$. If $\mathcal{D}$ and $\mathcal{D}'$ are in $Y_{\mathcal{D}_2}$, then they intersect.*

Proof. The proof is illustrated in Figure 3.6. Let p_1, p_3, p'_1, p'_3 be points in $\mathcal{D} \cap \mathcal{D}_1$, $\mathcal{D} \cap \mathcal{D}_3$, $\mathcal{D}' \cap \mathcal{D}_1$ and $\mathcal{D}' \cap \mathcal{D}_3$, respectively. Let us assume for a contradiction that $\mathcal{D}$ and $\mathcal{D}'$ do not intersect. Let p_2 be in $\mathcal{D} \cap \mathcal{D}_2$ and p'_2 be in $\mathcal{D}' \cap \mathcal{D}_2$. We claim that p_2 is not in the quadrilateral $p_1p_3p'_3p'_1$. By assumption, it is possible to continuously move the points p_1, p_3 in $\mathcal{D}_1$ and $\mathcal{D}_3$ respectively such that p_1 and p'_1 overlap, the points p_3 and p'_3 overlap, while keeping the property that

$[p_1, p_3]$ does not intersect $\mathcal{D}_2$. Observe that this is not possible if p_2 is in the quadrilateral $p_1p_3p_3'p_1'$, because $\mathcal{D}_1$ and $\mathcal{D}_3$ do not intersect with $\mathcal{D}_2$. Moreover by assumption $\mathcal{D}_2$ does not intersect with the segments $[p_1, p_3]$, $[p_3, p_3']$, $[p_3', p_1']$ and $[p_1', p_1]$. This implies that $\mathcal{D}_2$ is outside of the quadrilateral $p_1p_3p_3'p_1'$. Let us consider a separating line ℓ of $\mathcal{D}$ and $\mathcal{D}'$. As ℓ intersects the segments $[p_1, p_1']$, $[p_2, p_2']$ and $[p_3, p_3']$, we observe that ℓ is a line transversal of $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$. Let us consider the intersection s of ℓ with the quadrilateral $p_1p_3p_3'p_1'$. This intersection s could a priori be the union of two segments if the quadrilateral $p_1p_3p_3'p_1'$ is not convex, however that is not possible since ℓ does not intersect $[p_1, p_3]$ and $[p_1', p_3']$. We have that s is a segment with one endpoint in $\mathcal{D}_1$ and the other in $\mathcal{D}_3$. This implies that the disk in the middle of ℓ is not $\mathcal{D}_2$.

If $\mathcal{D}_2$ is the disk in the middle of all line transversals, we have already reached a contradiction. Let us now assume that at least one of $\mathcal{D}_1$ and $\mathcal{D}_3$ is 1-intersecting. Without loss of generality, let us assume that $\mathcal{D}_1$ is the disk in the middle of ℓ. We denote by τ_1 the external tangent of $\mathcal{D}_2$ and $\mathcal{D}_3$ which intersects $\mathcal{D}_1$. Both segments $[p_2, p_3] \cap \mathcal{D}_1$ and $[p_2', p_3'] \cap \mathcal{D}_1$ are on the same side of τ_1. By assumption, only one of the two segments can be continuously moved to τ_1 without touching ℓ. Without loss of generality, let us assume that this segment is $[p_2, p_3] \cap \mathcal{D}_1$. Observe that $[p_2, p_3]$ and p_1' are on different sides of ℓ. The line segment $[p_2, p_3]$ splits $\mathcal{D}_1$ into two parts, one of them being contained in $\mathcal{D}$. This implies that $\mathcal{D}$ contains p_1' or is centred with respect to $\mathcal{D}_1$, which is a contradiction. □

Lemma 3.19. *Let us assume that $\mathcal{D}_1$ is the disk in the middle of no line transversal or is 1-intersecting, and that the same holds with $\mathcal{D}_3$. If $\mathcal{D}$ is in $X_{\mathcal{D}_2}$ and $\mathcal{D}'$ is in $Y_{\mathcal{D}_2}$, then they intersect.*

Proof. Let p_1, p_3 be points coming from the definition of $\mathcal{D}$ being in X. Let p_1', p_3' be points in $\mathcal{D}' \cap \mathcal{D}_1$ and $\mathcal{D}' \cap \mathcal{D}_3$, respectively. By definition, the segment $[p_1, p_3]$ splits $\mathcal{D}_2$ into two parts. We denote by A_2 the part of $\mathcal{D}_2$ that is contained in $\mathcal{D}$. Assume for a contradiction that $\mathcal{D}$ and $\mathcal{D}'$ do not intersect. We claim that A_2 is inside the quadrilateral $p_1p_3p_3'p_1'$. Without loss of generality, let us assume that $[p_1', p_3'] \cap \mathcal{D}_1 = \{p_1'\}$ and $[p_1', p_3'] \cap \mathcal{D}_3 = \{p_3'\}$. It is possible to move p_1' and p_3' to τ, while following the boundaries of $\mathcal{D}_1$ and $\mathcal{D}_3$ respectively, such that $[p_1', p_3']$ does not intersect $\mathcal{D}_2$. Since $\mathcal{D}$ and $\mathcal{D}'$ do not intersect, it implies that $[p_1, p_3]$ and $[p_1', p_3']$ do not intersect. In particular, it implies that p_1' is closer to τ than p_1 (when considering them as points on the boundary of $\mathcal{D}_1$), and likewise p_3' is closer to τ than p_3 is. By assumption, $\mathcal{D}_2$ does not intersect the segments $[p_1, p_1']$, $[p_1', p_3']$ and $[p_3', p_3]$. We have shown that A_2 is inside the quadrilateral $p_1p_3p_3'p_1'$.

As $\mathcal{D}$ and $\mathcal{D}'$ are not intersecting, we have that $\mathcal{D}' \cap A_2$ is empty. Let us denote by ℓ a separating line of $\mathcal{D}$ and $\mathcal{D}'$. As ℓ does not intersect $[p_1, p_3]$ or $[p_1', p_3']$, but because ℓ intersects $[p_1, p_1']$ and $[p_3, p_3']$, we have that ℓ splits the quadrilateral $p_1p_3p_3'p_1'$. Furthermore, ℓ does not intersect A_2, and thus it is a line transversal of $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$ whose disk in the middle is not $\mathcal{D}_2$. If $\mathcal{D}_2$ is the disk in the middle of all line transversals, we have already reached a contradiction. Let us now assume that at least one of $\mathcal{D}_1$ and $\mathcal{D}_3$ is 1-intersecting. Without loss of generality, let us assume that the disk in the middle of ℓ is $\mathcal{D}_1$. Let τ_1 be the external tangent of $\mathcal{D}_2$ and $\mathcal{D}_3$ which intersects $\mathcal{D}_1$. Let p_2' be in $\mathcal{D}' \cap D_2$. Now, the segment $[p_2', p_3'] \cap \mathcal{D}_1$ lies between the lines τ_1 and ℓ, and $[p_2', p_3']$ splits $\mathcal{D}_1$ in such a way that $\mathcal{D}'$ either contains p_1 or is centred with respect to $\mathcal{D}_1$, which is a contradiction. □

We have now shown that under certain conditions, $G(X \cup Y)$ and $G(X' \cup Y')$ are complete graphs. We now prove three lemmas to show that $G(X \cup Y \cup Z)$ and $G(X' \cup Y' \cup Z')$ are complete graphs. To do so, we have to show that all pairs of disks in Z intersect. In the following lemma, we prove the stronger statement that all pairs of disks in $Z \cup Z'$ intersect. Then, in Lemmas 3.21 and 3.22, we show that a disk $\mathcal{D}$ in Z intersects any disk in X or Y. This implies that

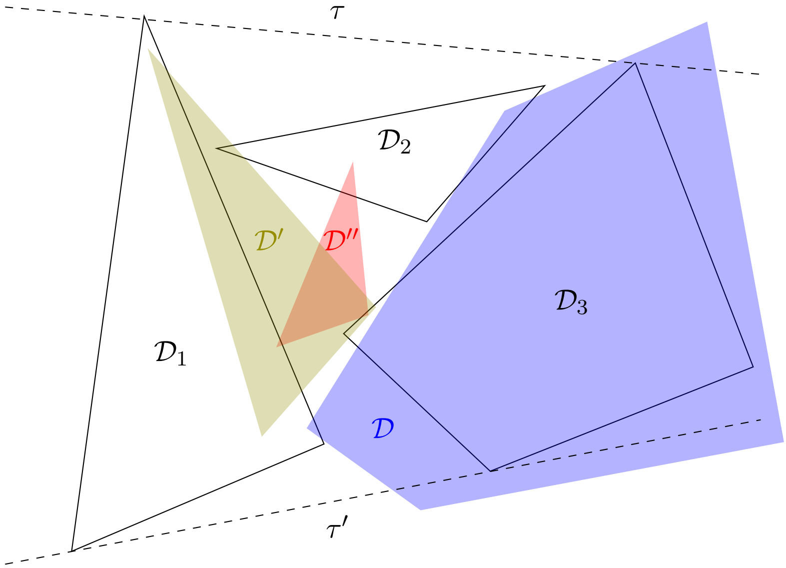

Figure 3.7: The disk $\mathcal{D}$ is in Z, the disk $\mathcal{D}'$ is in X' and the disk $\mathcal{D}''$ is in Y'. Observe that $\mathcal{D}$ does not intersect $\mathcal{D}'$ or $\mathcal{D}''$.

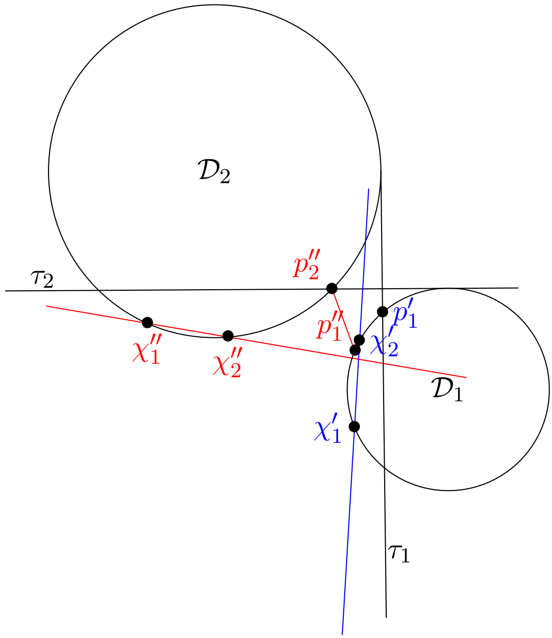

Figure 3.8: Illustration of the proof of Lemma 3.20.

if $\mathcal{D}$ does not intersect a disk $\mathcal{D}'$ in $\{\mathcal{D}'_j\}$, then $\mathcal{D}'$ is in $X' \cup Y'$, as illustrated in Figure 3.7.

Lemma 3.20. *Let $\mathcal{D}'$ and $\mathcal{D}''$ be intersecting with $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$. If $\mathcal{D}'$ and $\mathcal{D}''$ are respectively centred with respect to $\mathcal{D}_i$ and $\mathcal{D}_j$, $i, j \in \{1, 2, 3\}$, then they intersect.*

Proof. The proof is illustrated in Figure 3.8. For a disk $\mathcal{D}_i$ in $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$ which is 1-intersecting, we denote by τ_i the external tangent of the two other disks that $\mathcal{D}_i$ intersects. Furthermore, we denote by A_i the subset of $\mathcal{D}_i$ which lies on the same side of τ_i as the two other disks. Let $\mathcal{D}'$ and $\mathcal{D}''$ be two disks in that are centred. If they both contain the same subset $\mathcal{D}_i \setminus A_i$ for some $i \in \{1, 2, 3\}$, then they intersect. Otherwise, let us assume without loss of generality that $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$ and $\mathcal{D}''$ is centred with respect to $\mathcal{D}_2$. There are two intersections between the boundaries of $\mathcal{D}''$ and $\mathcal{D}_2$, that we denote by χ_1'' and χ_2''. We denote by $\mathcal{H}''$ the closed halfplane with bounding line (χ_1'', χ_2'') that contains $\mathcal{D}_2 \setminus A_2$. We denote by $\mathcal{H}_2$ the closed halfplane with bounding line τ_2 that contains A_2. By assumption, $\mathcal{D}_1$ intersects only one of the two connected sets of $(\mathcal{H}'' \cap \mathcal{H}_2) \setminus A_2$. Let us consider the intersections of τ_2 with the boundary of $\mathcal{D}_2$. By what we just said, there is a closest intersection to $\mathcal{D}_1$, that we denote by p_2''. Note that p_2'' is in $\mathcal{H}''$, and therefore in $\mathcal{D}''$. Let p_1'' be a point in $\mathcal{D}'' \cap D_1$. If p_1'' is in $\mathcal{D}'$ then we are done. Let us now assume that it is not the case. We denote by χ_1' and χ_2' the intersections of $\mathcal{D}'$ with the boundary of $\mathcal{D}_1$. Without loss of generality, we can assume that p_1'' is on the boundary of $\mathcal{D}_1$. We denote by p_1' the intersection of τ_1 and the boundary of $\mathcal{D}_1$ that is the closest to $\mathcal{D}_2$, which can be defined similarly to how we defined p_2''. Now observe that one of χ_1' and χ_2' is between p_1'' and p_1' on the boundary of $\mathcal{D}_1$. Assume without loss of generality that χ_2' is the closest to $\mathcal{D}_2$. Let us consider the halfplane $\mathcal{H}'$ with bounding line (χ_1', χ_2') that contains $\mathcal{D}_1 \setminus A_1$. We also denote by $\mathcal{H}_1$ the halfplane with bounding line τ_1 that contains A_1. As $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$, there is one of the two connected component that intersects with $\mathcal{D}_2$, and the other with $\mathcal{D}_3$. Note that the connected component on the side of χ_2' cannot intersect with $\mathcal{D}_3$, since otherwise $\mathcal{D}_1$ would not be the disk in the middle of τ_1. This implies that the connected component on the side of χ_2' is the one that intersects $\mathcal{D}_2$. Finally, observe that either $\mathcal{D}'$ contains p_2'', or it does not intersect with A_2, and thus contains a point in $\mathcal{D}_2 \setminus A_2$. In both cases, $\mathcal{D}'$ contains a point in $\mathcal{D}''$. □

Lemma 3.21. *Let us assume that $\mathcal{D}_1$ or $\mathcal{D}_3$ is 1-intersecting. If $\mathcal{D}$ is in $X_{\mathcal{D}_2}$ and $\mathcal{D}'$ is in $Z_{\mathcal{D}_2}$, then they intersect.*

Proof. Let p_1, p_3 be the points coming from the definition of $\mathcal{D}$ being in X. By assumption, the segment $[p_1, p_3]$ splits $\mathcal{D}_2$ into two parts. Let us denote by A_2 the part that is contained in $\mathcal{D}$. Without loss of generality, let us assume that $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$. We have that p_1 is on the boundary of $\mathcal{D}_1$. Let us denote by p'_1 and q'_1 the two intersections between the boundaries of $\mathcal{D}'$ and $\mathcal{D}_1$. Let us consider the boundary of $\mathcal{D}_1$. The lines τ and τ' cut it into two parts, one of them containing p_1, p'_1 and q'_1. Let us consider that part of the boundary of $\mathcal{D}_1$. We assume that when going from τ to τ' while following this part, we see the points in that order: p'_1, p_1 and then q'_1. Indeed if p_1 does not appear between p'_1 and q'_1, then p_1 is in $\mathcal{D}'$ and we are done. We can assume without loss of generality that we reach first p'_1 and last q'_1 by relabelling if need be. Let us follow the boundary of $\mathcal{D}'$ from p'_1 while staying outside of $\mathcal{D}_1$ (therefore not going in the direction of q'_1), until we reach either $\mathcal{D}_2$ or $\mathcal{D}_3$. This must happen as $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$. We claim that we can only reach $\mathcal{D}_2$. Indeed, assume for a contradiction that we reach a point p'_3 at an intersection between the boundaries of $\mathcal{D}'$ and $\mathcal{D}_3$. As $\mathcal{D}'$ is in Z, the segment $[p'_1, p'_3]$ does not intersect $\mathcal{D}_2$. However, the segment $[p_1, p_3]$ does intersect $\mathcal{D}_2$. As $[p_1, p_3]$ and $[p'_1, p'_3]$ do not intersect, and since p'_1 is closer to τ than p_1 is (when considering their positions on the boundary of $\mathcal{D}_1$), it implies that p'_3 is closer to τ than p_3 is (when considering their positions on the boundary of $\mathcal{D}_3$). It is now possible to continuously move p'_1 and p'_3 to τ while keeping the property that $[p'_1, p'_3]$ does not intersect $\mathcal{D}_2$, which is impossible. We have shown that when following the boundary of $\mathcal{D}'$ from p'_1, we meet a point p'_2 in $\mathcal{D}' \cap \mathcal{D}_2$. By construction, as p'_1 is closer to τ than p_1 is, we know that p'_2 is in A_2, which implies that $\mathcal{D}$ and $\mathcal{D}'$ intersect. □

Lemma 3.22. *Let us assume that $\mathcal{D}_1$ or $\mathcal{D}_3$ is 1-intersecting. If $\mathcal{D}$ is in $Y_{\mathcal{D}_2}$ and $\mathcal{D}'$ is in $Z_{\mathcal{D}_2}$, then they intersect.*

Proof. Assume for a contradiction that they do not intersect. Let p_1 and p_3 be in $\mathcal{D} \cap \mathcal{D}_1$ and $\mathcal{D} \cap \mathcal{D}_3$ respectively. We can even assume that $[p_1, p_3] \cap \mathcal{D}_1 = \{p_1\}$ and $[p_1, p_3] \cap \mathcal{D}_3 = \{p_3\}$. Without loss of generality, let us assume that $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$. We define p'_1 and q'_1 as in the proof of Lemma 3.21. We follow the boundary of $\mathcal{D}'$ from p'_1 while staying outside of $\mathcal{D}_1$ (therefore not going in the direction of q'_1), until we reach either $\mathcal{D}_2$ or $\mathcal{D}_3$. Again, this must happen as $\mathcal{D}'$ is centred with respect to $\mathcal{D}_1$. However, we cannot reach $\mathcal{D}_3$ because if so, by denoting p'_3 the point on the boundary of $\mathcal{D}_3$ that we reach, the points p'_1 and p'_3 could be moved to τ while having $[p'_1, p'_3]$ not intersecting $\mathcal{D}_2$. This is because p'_1 is closer to τ than p_1 is (when considering their positions on the boundary of $\mathcal{D}_1$), and the same holds with p'_3 and p_3. Since $\mathcal{D}$ is in Y, $[p_1, p_3]$ can be moved to τ without intersecting $\mathcal{D}_2$, and thus so can $[p'_1, p'_3]$. This implies that the disk we reach is $\mathcal{D}_2$. But this is in contradiction with the fact that $[p_1, p_3]$ can be moved to τ without intersecting $\mathcal{D}_2$, since p'_1 is closer to τ than p_1 is. □

We can now prove Theorem 3.2.

Proof of Theorem 3.2. We consider any fixed representation of G with convex pseudo-disks. We denote by $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$ the three non-intersecting sets corresponding to H. Likewise we denote by $\{\mathcal{D}'_j\}$ the convex pseudo-disks in $G \setminus H$. If there is no line transversal of $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$, we conclude with Lemma 3.9. If there is one disk that is 2-intersecting, say $\mathcal{D}_2$, then we have that Y, Y', Z and Z' are empty. We conclude with Lemma 3.17. Now let us assume that one disk, say $\mathcal{D}_2$ is contained. Then we know with Lemma 3.13 that $\mathcal{D}_1$ is the disk in the middle of no line transversal, or it 1-intersecting. The same holds with $\mathcal{D}_3$. Therefore we can apply Lemmas 3.17, 3.18, 3.19, 3.20, 3.21 and 3.22. They imply that $G(X \cup Y \cup Z)$ is a complete graph. By the same arguments, $G(X' \cup Y' \cup Z')$ is a complete graph too. As $\{\mathcal{D}'_j\}$ is the disjoint union of X, X', Y, Y', Z and Z', it implies that $G \setminus H$ is cobipartite. If no disk is 2-intersecting and no disk is contained, then all disks are either the disk in the middle of no line transversal, or are 1-

intersecting. As we are now assuming that there is a line transversal of $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$, we assume without loss of generality that its disk in the middle is $\mathcal{D}_2$, and we use Lemmas 3.17, 3.18, 3.19, 3.20, 3.21 and 3.22 to conclude. □

3.4 3-pancakes and disks

In this section, we show the following theorem:

Theorem 3.23. *If Conjecture 3.4 holds, there exists a randomised EPTAS for computing a maximum clique in Π^3, even without a representation.*

We first give some definitions. Vapnik and Chervonenkis have introduced the concept of VC-dimension in [81]. In this paper, we are only concerned with the VC-dimension of the neighbourhood of some geometric intersection graphs. In this context, the definition can be stated as follows:

Definition 3.24. *Let $\mathcal{F}$ be a family of sets in $\mathbb{R}^d$, and let G be the intersection graph of $\mathcal{F}$. We say that $F \subseteq \mathcal{F}$ is* shattered *if for every subset X of F, there exists a vertex v in G that is adjacent to all vertices in X, and adjacent to no vertex in $F \setminus X$. The VC-dimension of the neighbourhood of G is the maximum cardinality of a shattered subset of $\mathcal{F}$.*

We define the class $\mathcal{X}(d, \beta, K)$ as introduced by Bonamy *et al.* in [8]. Let d and K be in $\mathbb{N}$, and let β be a real number such that $0 < \beta \leq 1$. Then $\mathcal{X}(d, \beta, K)$ denotes the class of simple graphs G such that the VC-dimension of the neighbourhood of G is at most d, $\alpha(G) \geq \beta|V(G)|$, and $\text{iocp}(G) \leq K$. They show that there exist EPTAS (Efficient Polynomial-Time Approximation Scheme) for computing a maximum independent set in $\mathcal{X}(d, \beta, K)$. An EPTAS for a maximisation problem is an approximation algorithm that takes a parameter $\varepsilon > 0$ and outputs a $(1-\varepsilon)$-approximation of an optimal solution, and running in $f(\varepsilon)n^{\mathcal{O}(1)}$ time. More formally, we have the following:

Theorem 3.25 (Bonamy *et al.* [8])**.** *For any constants $d, K \in \mathbb{N}$, $0 < \beta \leq 1$, for every $\varepsilon > 0$, there is a randomised $(1-\varepsilon)$-approximation algorithm running in time $2^{\tilde{\mathcal{O}}(1/\varepsilon^3)}n^{\mathcal{O}(1)}$ for maximum independent set on graphs of $\mathcal{X}(d,\beta,K)$ with n vertices.*

Recently, Dvořák and Pekárek have announced that it is not necessary to have bounded VC-dimension [27]. More explicitly, there is an EPTAS for the class $\mathcal{X}(+\infty,\beta,K)$. However, their running time dependence in n is higher: $\tilde{\mathcal{O}}(n^5)$ with Dvořák and Pekárek's algorithm compared to $\tilde{\mathcal{O}}(n^2)$ with the one of Bonamy *et al.* Also, Dvořák and Pekárek do not compute the dependence in ε. For this reason, we prefer the algorithm of Bonamy *et al.*, despite the fact that we have to show bounded VC-dimension.

Theorem 3.25 states that there exists an EPTAS for computing a maximum independent set on graphs of $\mathcal{X}(d,\beta,K)$, for any $d, K \in \mathbb{N}$ and $0 < \beta \leq 1$. Let G be in Π^3. In order to prove Theorem 3.23, we show that the VC-dimension of the neighbourhood of any vertex in G is bounded. Observe that the VC-dimension of a graph and its complement are equal. We aim at using the EPTAS mentioned above for computing a maximum independent set in the complement, which is equivalent to computing a maximum clique in the original graph. However a graph G in Π^3 does not necessarily satisfy $\alpha(\overline{G}) \geq \beta|V(G)|$ for some $0 < \beta \leq 1$. Even if it does, we need to know the value of β in order to use the EPTAS of Theorem 3.25. Therefore we show how to compute a maximum clique in any $G \in \Pi^3$ by using polynomially many times the EPTAS of Theorem 3.25 on some subgraphs of G, which have the desired property.

In general, for intersection graphs of geometric objects that can be described with finitely many parameters, the VC-dimension of the neighbourhood is bounded. For graphs in Π^3, we were able to show an upper bound of 28. We do not expect this value to be tight, but showing any constant was sufficient for our purpose.

Proposition 3.26. *The VC-dimension of the neighbourhood of a graph $G = (V, E)$ in Π^3 is at most* 28*.*

We use the fact that the VC-dimension of the neighbourhood of disk graphs (and even pseudo-disk graphs) is at most 4, as proved by Aronov *et al.* [4]. Likewise, the VC-dimension of the neighbourhood of unit ball graphs is at most 4, as noticed by Bonamy *et al.* [8]. For any point $c \in \mathbb{R}^3$ and any non-negative real number ρ, we denote by $\mathcal{B}(c, \rho)$ the ball centred at c with radius ρ. Moreover, we denote by $P^3(c, \rho)$ the 3-pancake that is the Minkowski sum of the unit ball centred at the origin and the disk lying on the plane xOy, centred at c with radius ρ. Note that if $\rho = 0$, then $P^3(c, \rho)$ is the unit ball centred at c. Before showing Proposition 3.26, we show the following:

Lemma 3.27. *Let $\mathcal{B}$ be a unit ball centred at c and let $P^3(c', \rho)$ be a 3-pancake. We denote by $\mathcal{D}$ the disk that is the intersection of $\mathcal{B}(c, 2)$ and the plane xOy. Also, we denote by $\mathcal{D}'$ the disk $\mathcal{D}(c', \rho)$ (which is a strict subset of the intersection of P^3 and the plane xOy). We have that $\mathcal{B}$ and P^3 intersect if and only if $\mathcal{D}$ and $\mathcal{D}'$ intersect.*

Proof. By definition, $\mathcal{B}$ and P^3 intersect if and only if there exists a unit ball $\mathcal{B}'$ whose centre lies in $\mathcal{D}'$ such that $\mathcal{B}$ and $\mathcal{B}'$ intersect. This is equivalent to say that $\mathcal{B}(c, 2)$ contains a point in $\mathcal{D}'$. Finally, this statement is equivalent to having $\mathcal{D}$ and $\mathcal{D}'$ intersecting. □

Proof of Proposition 3.26. First let us show that if V is shattered, then in any Π^3 representation of G there are at most four 3-pancakes. Let us assume by contradiction that there exists a set S of five 3-pancakes, such that for every subset T of S, there exists a unit ball or a 3-pancake intersecting all elements in T and intersecting no element in $S \setminus T$. For each 3-pancake $P^3(c_i, \rho_i)$ in S, we denote by $\mathcal{D}_i$ the disk $\mathcal{D}(c_i, \rho_i)$ lying on the plane xOy. Let T be a subset of S. If there exists a 3-pancake $P^3(c', \rho')$ intersecting with the elements of T and with no element in $S \setminus T$, we denote by $\mathcal{D}_T$ the disk $\mathcal{D}(c', \rho' + 2)$ lying on the plane xOy. Otherwise there exists a unit ball $\mathcal{B}$ centred at c'' intersecting intersecting with the elements of T and with no element in $S \setminus T$, and then we denote by $\mathcal{D}_T$ the intersection between $\mathcal{B}(c'', 2)$ and xOy. As $\mathcal{B}$ intersects with a 3-pancake, $\mathcal{D}_T$ is not empty. Using Lemma 3.27, we have that $\mathcal{D}_i$ intersects with $\mathcal{D}_T$ if and only if $P^3(c_i, \rho_i)$ is in T. This implies that if S is shattered by some 3-pancakes and unit balls, then the set $\{\mathcal{D}_i\}$ is shattered by $\{\mathcal{D}_T \mid T \subseteq S\}$. However this is not possible because the VC-dimension of the neighbourhood of disk graphs is at most 4.

Now let us prove the claim. Assume by contradiction that we have a shattered set with 29 elements. As shown above, in any Π^3 representation there are at least 25 unit balls. Let us consider such a representation. We denote by $S_1, \ldots, S_5$ five sets of five unit balls each. As the VC-dimension of the neighbourhood of unit ball graphs is at most 4, for each set S_i there exists a non-empty subset $T_i \subseteq S_i$ such that no unit ball can intersect with the unit balls in T_i, but not with those in $S_i \setminus T_i$. Therefore the absolute height of the centre of any unit ball in T_i is at most 2, since T_i is realised by a 3-pancake. For each T_i, we choose arbitrarily one unit ball $\mathcal{B}_i$, and define a new set T as $\{\mathcal{B}_1, \ldots, \mathcal{B}_5\}$. Moreover for each unit ball $\mathcal{B}_i$ centred at c_i, we denote by $\mathcal{D}_i$ the intersection between $\mathcal{B}(c_i, 2)$ and the plane xOy. Note that $\mathcal{D}_i$ is not empty. Let T' be a subset of T, and let us consider the set $\cup_{\mathcal{B}_i \in T'} T_i$, that we denote by T'_+. Note that unless $T' = \emptyset$, no unit ball can intersect with all elements in T'_+ and with no element in $S \setminus T'_+$. Therefore this can only be achieved by a 3-pancake $P^3(c, \rho)$, and we denote by $\mathcal{D}_{T'}$ the disk $\mathcal{D}(c, \rho)$ lying on the plane xOy. Using Lemma 3.27, the five disks $\mathcal{D}_i$ are shattered by the disks in $\{D_{T'} \mid T' \subseteq T\}$, which is impossible. □

Proof of Theorem 3.23. Let G be a graph in Π^3 with n vertices. Since the VC-dimension of a graph is the same as its complement, Proposition 3.26 implies that the VC-dimension of $\overline{G}$ is at most 28. First let us assume that a representation of G is given. For every vertex represented by a unit ball, we are going to compute a maximum clique containing this vertex. As noticed by Bonamy *et al.*, for any vertex v represented by a unit ball, we have $|\mathcal{N}(v)| \leq 25\omega(G)$ [8]. Let us denote by G_v the subgraph induced by $\mathcal{N}(v)$. Thus we have $\alpha(\overline{G}_v) \geq |\mathcal{N}(v)|/25$. This shows that $\overline{G}_v$ is in $\mathcal{X}(28, 1/25, K)$. Using Theorem 3.25, we have a randomised EPTAS for computing a maximum independent set in $\overline{G}_v$, which is equivalent to computing a maximum clique in G_v. Note that computing a maximum clique in G_v for each vertex v represented by a unit ball adds at most a multiplicative factor n in the running time. It remains to compute a maximum clique that only contains vertices represented by 3-pancakes. Instead of considering 3-pancakes, one can only look at the corresponding disks on the plane xOy. This can be done as suggested in [8]: find four piercing points in time $\mathcal{O}(n^8)$, then consider the subgraph H of disks that are pierced by at least one of these points. We have $\alpha(\overline{H}) \geq n'/4$ where n' denotes the number of vertices in H. This implies that H is in $\mathcal{X}(28, 1/4, K)$, and we can conclude as before.

Now assume that a representation is not given. As we do not know whether a vertex can be represented by a unit ball, we cannot compute a maximum clique as was done above. If there exists a representation of G with at least one vertex v represented as a unit ball, then $\alpha(G_v) \leq 12$, because the kissing number for unit spheres is 12. Indeed for any 3-pancake P^3 intersecting a unit ball B, there exists a unit ball $B' \subseteq P^3$ such that B and B' intersect. Thus, if instead of each pancake there were such a unit ball, we would have the desired inequality. But since such a unit ball B' is contained in the corresponding 3-pancake P^3, the independence number of G_v can only decrease when considering the actual 3-pancakes, which implies $\alpha(G_v) \leq 12$. If there exists a representation only with 3-pancakes, then the vertex v corresponding the 3-pancake with the smallest radius satisfies $\alpha(G_v) \leq 6$. Therefore in any case there must be a vertex v with $\alpha(G_v) \leq 12$. We can find such a vertex in $\mathcal{O}(n^{13})$ time by testing for each v whether there is an independent of size 12 in G_v.

In order to give a linear lower bound on $\alpha(\overline{G}_v)$, we first give an upper bound on the chromatic number of any graph in Π^3. Let $\tilde{G}$ be a graph in Π^3, given with a fixed representation. We denote by V_1 the set of vertices represented by unit balls, and by V_2 those represented by 3-pancakes. We denote by $\tilde{G}_1$ the graph induced by V_1. As noted in [8], we have for each $v_1 \in V_1$, $|\mathcal{N}(v_1)| \leq 25\omega(\tilde{G}_1)$. Since $\omega(\tilde{G}_1) \leq \omega(\tilde{G})$, the maximum degree in $\tilde{G}_1$ is at most $25\omega(\tilde{G}) - 1$, which implies that we can colour the vertices in V_1 using at most $25\omega(\tilde{G})$ colours. For disk graphs, the chromatic number is at most 6 times the clique number. Thus we can colour the vertices in V_2 using at most $6\omega(\tilde{G})$ other colours. So in total we have $\chi(\tilde{G}) \leq 31\omega(\tilde{G})$.

Let us consider again the subgraph G_v, we have $\alpha(G_v)\omega(G_v) \geq \alpha(G_v)\chi(G_v)/31 \geq |\mathcal{N}(v)|/31$. Therefore, we obtain the inequality $\omega(G_v) \geq |\mathcal{N}(v)|/372$. This implies that $\overline{G}_v$ is in $\mathcal{X}(28, 1/372, K)$, and therefore we have an EPTAS for computing a maximum clique containing v. We can iterate this process in the graph G where v has been removed to compute a maximum clique that does not contain v. As we repeat this process linearly many times, we obtain an EPTAS for computing a maximum clique in G. □

CHAPTER 4

Minimum convex partition

4.1 Introduction

In Chapters 2 and 3, we discussed the maximum clique problem in intersection graphs (of disk-related shapes). In Chapters 4 and 5, we study problems on geometric plane graphs on point sets in degenerate position. The two parts of this thesis are morally independent.

The CG Challenge 2020 organised by Demaine, Fekete, Keldenich, Krupke and Mitchell [24], was about solving instances of *Minimum Convex Partition* (MCP).

Definition 4.1 (Demaine *et al.* [24]: Minimum Convex Partition problem)**.** Given a set P of n points in the plane. The objective is to compute a plane graph with vertex set P (with each point in P having positive degree) that partitions the convex hull of P into the smallest possible number of convex faces. Note that collinear points are allowed on face boundaries, so all internal angles of a face are at most π.

As explained by Bose *et al.*, this problem has applications in routing [12]. They showed that a routing algorithm named *Random-Compass* that works for triangulations can be extended to convex partitions. Having a convex partition with few faces reduces the amount of data to store. From now on, we denote by P a set of n points in the plane.

In this chapter, we present several approximation algorithms for MCP. We obtain those approximation algorithms by relating the MCP problem to the *Covering Points with Non-Crossing Segments* (CPNCS) problem. First, we define what *non-crossing segments* are.

Definition 4.2 (Non-Crossing Segments)**.** We call a part of a line bounded by two points a *segment*. The two points are referred to as *endpoints* of the segment. Note that we do not force the endpoints to be distinct, therefore we consider a point p as being a segment. The endpoint of p is p itself. Two segments are *non-crossing* if the intersection of their relative interior is empty.

Definition 4.3 (Covering Points with Non-Crossing Segments)**.** Let P be a set of n points. Find a minimum number of non-crossing segments whose endpoints are in P such that each point of P is contained in at least one segment.

The condition that the endpoints of the segments must be in P has no effect on the number of segments required. We add it as it simplifies some arguments. Note that CPNCS is not a so-called *set* *cover problem* nor an *exact cover problem*. We believe that CPNCS is interesting in itself. Even though it is a very natural problem, to the best of our knowledge it had not been introduced before.

4.1.1 NP-hardness results

Fevens, Meijer and Rappaport first considered the MCP problem in 2001 [29], and its complexity was explicitly asked about by Knauer and Spillner in 2006 [52]. It has remained open since then [6, 24]. We show in Section 4.4 that MCP is NP-hard. To do this, we use the decision version of the problem, as stated below:

Definition 4.4 (MCP - decision version)**.** Given a set P of points in the plane and a natural number k, is it possible to find at most k closed convex polygons whose vertices are points of P, with the following properties:

- The union of the polygons is the convex hull of P,
- The interiors of the polygons are pairwise disjoint,
- No polygon contains a point of P in its interior.

We also show NP-hardness of a similar problem, which we call *Minimum Convex Tiling* problem (MCT). The problem is exactly as in Definition 4.4, but the constraint about the vertices of the polygons is removed (i.e. they need not be points of P). This can make a difference as shown in Figure 4.1. Equivalently, the MCT problem corresponds to the MCP problem when Steiner points are allowed. A *Steiner point* is a point that does not belong to the point set given as input, and which can be used as a vertex of some polygons. The MCT problem has been studied in 2012 by Dumitrescu, Har-Peled and Tóth, who asked about the complexity of the problem [25]. We answer their question, and our proofs are very similar for MCP and MCT.

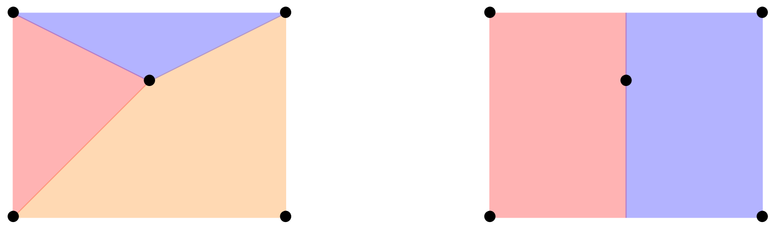

Figure 4.1: A minimum partition with three convex polygons and a tiling with two.

We show in Section 4.4.5 that CPNCS is NP-hard, even for some constrained point sets, using a reduction from *Maximum Independent Set in Intersection Graphs of Segments.*

4.1.2 Approximation algorithms

For the related problem *Minimum Convex Partition of Polygons with Holes*, Bandyapadhyay, Bhowmick and Varadarajan showed the existence of a $(1+\varepsilon)$-approximation algorithm running in time $n^{\mathcal{O}((\log n/\varepsilon)^4)}$ [5]. Although they only consider holes with non empty interior, one can observe that their proof extends to the case of point holes. This is an even more general setting than MCP for point sets, so their algorithm also applies in our setting. This implies that MCP is not APX-hard unless $NP \subseteq DTIME(2^{\text{polylog } n})$.

Under the assumption that no three points are collinear, Knauer and Spillner have shown a $\frac{30}{11}$-approximation algorithm [52] for MCP in 2006. As a lower bound on the number of convex faces for one particular point set, they rely on the observation that each inner point has degree at least 3. This gives a lower bound on the number of edges, and therefore on the number of faces, by Euler's formula. Note that the restriction that no three points are on a line is necessary, as shown in Figure 4.2. There are only two faces in a minimum convex partition of this point set, and all the inner points have degree 2.

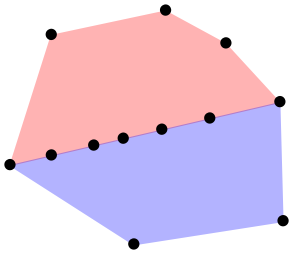

Figure 4.2: The number of inner points can be arbitrarily much larger than the number of convex faces required.

Additionally, Knauer and Spillner showed how to adapt any constructive upper bound on the number of faces into an approximation algorithm. More explicitly, they showed that if one can compute in polynomial time a convex partition with at most λn convex faces, then there exists a 2λ-approximation algorithm running in polynomial time. The best result to date is a proof by Sakai and Urrutia that one can partition a point set in quadratic time using at most $\frac{4}{3}n$ convex faces (the result was presented at the 7th JCCGG in 2009, the paper appeared on arXiv in 2019) [71]. Although they do not mention it, combining this result with the one by Knauer and Spillner gives a quadratic time $\frac{8}{3}$-approximation algorithm.

Concerning previous upper bounds, Neumann-Lara, Rivero-Campo and Urrutia first showed in 2004 how to construct in quadratic time a partition of any point set with at most $\frac{10}{7}n$ convex faces [65]. In 2006, Knauer and Spillner improved this to $\frac{15}{11}n$ convex faces [52]. As said above, the best known upper bound is $\frac{4}{3}n$, as proven by Sakai and Urrutia in 2009.

Relatedly for lower bounds, García-Lopez and Nicolás have given in 2013 a construction of point sets for which any convex partition has at least $\frac{35}{32}n - \frac{3}{2}$ faces [31].

All these results concerning upper bounds hold for all point sets, even where many points are on a line. Indeed, slightly shifting the points so that no three points are on a line can only increase the number of convex faces needed. So an upper bound for point sets where no three points are on a line also holds for all point sets. However, as mentioned above, the lower bound used by Knauer and Spillner does not extend to our setting, where we consider all point sets. They say that a constant-approximation algorithm would be desirable for unrestricted point sets, but so far not even an $\mathcal{O}(n^{1-\varepsilon})$-approximation is known. For the MCT problem, Dumitrescu, Har-Peled and Tóth showed the existence of a 3-approximation algorithm for point sets with no three collinear points [25]. They also ask whether a constant-approximation algorithm exists when this constraint is removed. However, so far no $\mathcal{O}(n^{1-\varepsilon})$-approximation algorithm is known. In Section 4.2.2, we prove the following:

Theorem 4.5. *There exist $\mathcal{O}(\log OPT)$-approximation algorithms for MCP, MCT and CPNCS running in $\mathcal{O}(n^8)$-time.*

Allowing several points to be on a line does not simply create tedious technicalities to deal with. The crux of the matter is to find, for a fixed point set, an exploitable lower bound on the number of faces in a minimum convex partition. When no three points are on a line, the number of inner points in P gives a linear lower bound on the number of faces in a convex partition [52], and in a convex tiling [25]. In this chapter, we consider point sets with no restriction. We introduce the CPNCS problem as it pinpoints where the difficulty of finding a constant-approximation algorithm for MCP is and makes the problem easier to study. The *inner points* of P are the points not on the boundary of the convex hull. We show in Section 4.2.1 the following:

Theorem 4.6. *Let P be a set of n points with at least one inner point, and let $\lambda \geq 1$ be a real number. Let f_m denote the minimum* *number of faces in a convex partition of P. Let s_m denote the minimum number of non-crossing segments in a covering of the inner points of P, denoted by P_i.*

1. *It holds that $\frac{s_m}{6} \leq f_m \leq 8s_m$.*
2. *Given a covering of P_i with $s \leq \lambda s_m$ non-crossing segments, it is possible to compute in $\mathcal{O}(n^2)$-time a convex partition of P with at most $24\lambda f_m$ convex faces.*
3. *Given a convex partition of P with $f \leq \lambda f_m$ convex faces, it is possible to compute in $\mathcal{O}(n)$-time a covering of P_i with at most $44\lambda s_m$ non-crossing segments.*

The theorem also holds when considering convex tilings instead of convex partitions.

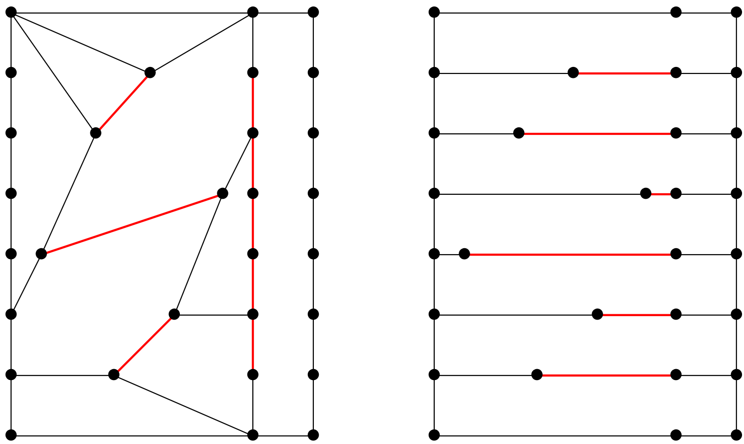

Figure 4.3: On the left side, a minimum covering of the inner points of P with 4 segments. A convex partition which contains those segments has at least 9 convex faces. On the right side, a covering of the inner points of P with 6 segments. There exists a minimum convex partition of P with 7 faces, which contains those segments.

Remark 4.7. The idea behind the similarity of MCP, MCT and CPNCS is that they are all about maximizing the number of vertices of degree 2 with incident edges being aligned in a plane straight-line drawing of a graph on a point set. We illustrate in Figure 4.3 that MCP and CPNCS are not strictly equivalent: We give a point set P in which the unique minimum convex partition of P does not contain the segments of any minimum covering of the inner points of P with non-crossing segments. Let us denote by P_s the set of endpoints of the three non-vertical segments of the covering. In any convex partition of P that contains those four segments of the covering, the points in P_s need to have degree at least 3. Moreover, each of them has to be connected to at least one point not in P_s. Finally, the topmost and bottommost of those six endpoints must be connected to at least two points not in P_s. Therefore, the convex partition drawn on the left is one that minimises the number of edges (and thus of faces) among the convex partitions that contain this minimum covering. This implies that finding a minimum covering of the inner points of some point set P with non-crossing segments does not necessarily help in finding a minimum convex partition of P. Nonetheless, Theorem 4.6 states that such a covering leads to an approximation for the MCP problem.

We call the algorithm for CPNCS the algorithm that iteratively picks a new segment among the valid ones that cover as many points not yet covered as possible, until all points in P are covered, the *greedy* algorithm. As we consider points to be potential segments, the algorithm terminates. We prove the following in Section 4.4.6.

Theorem 4.8. *There exist point sets for which the greedy algorithm for solving CPNCS realises an $\Omega(\sqrt{n})$-approximation.*

The CPNCS problem bears a resemblance with *Covering Points with Lines* (CPL), defined below, that we use in one of our approximation algorithms.

Definition 4.9 (Covering Points with Lines)**.** Given a set P of n points, find a minimum number of lines such that each point of P is contained in at least one line.

Before going into the proofs, we want to make a remark that we deem interesting. In [54], Anil Kumar, Arya and Ramesh mention that the CPL problem was motivated by the problem of *Covering a Rectilinear Polygon with Holes using Rectangles.* They say that getting a $o(\log n)$-approximation for this problem seems to require a better understanding of CPL. However, they are "not sure of the exact nature of this relationship". In this chapter, we show the hardness of MCP by using tools developed by Lingas to show NP-hardness of *Minimum Rectangular Partition for Rectilinear Polygons with Holes* [58]. The difference between a covering and a partition is that, in the latter, objects are interior-disjoint. Moreover, we prove that obtaining a constant-approximation algorithm for MCP is equivalent to finding one for CPNCS. Again, CPNCS is the non-crossing version of CPL. We hope our results help to better understand the relationship between these problems.

4.1.3 Exact algorithms, FPT algorithms

Under the assumptions that the points lie on the boundaries of a fixed number h of nested convex hulls, and that no three points lie on a line, Fevens, Meijer and Rappaport gave an algorithm for solving MCP in time $\mathcal{O}(n^{3h+3})$ [29]. Observe that this is not an FPT algorithm. Some integer linear programming formulations of the problem have been recently introduced [6, 74, 17].

A first FPT algorithm with respect to the number k of inner points was introduced by Grantson and Levcopoulos, with running time $\mathcal{O}(2^{16k}k^{6k-5}n)$ [33]. The idea of the algorithm is to enumerate all plane graphs on the inner points, and then for each to them to guess how to connect the inner points to the points on the boundary of the convex hull. Another FPT algorithm with respect to the number of inner points was later found by Spillner, with running time $\mathcal{O}(2^k k^4 n^3 + n \log n)$ [75].

We show in Section 4.3 the existence of an FPT algorithm that checks whether there is a solution for CPNCS with at most k non-crossing segments, running in time $\mathcal{O}(2^{k^2} k^{7k} + n^4 \log n)$. By Theorem 4.6, this gives us a constant-approximation FPT algorithm for MCP and MCT, where the parameter is the number of convex faces needed. Under the assumption that no three points are on a line, the number of faces in a minimum convex partition or in a minimum convex tiling is the same as the number of inner points, up to a constant multiplicative factor [52, 25]. However, when removing this assumption, the number of inner points can be arbitrarily much larger than the minimum number of convex faces, as shown in Figure 4.2. Our algorithm runs in time $\mathcal{O}(2^{36f^2} f^{42f+1} + n^4 \log n)$, where f denotes the minimum number of convex faces needed in a convex partition or in a convex tiling.

4.2 Approximation algorithms

4.2.1 The relation between MCP, MCT and CPNCS

Throughout this section, we denote by P a point set in the plane. We denote by P_i the set of inner points of P. Let p be in P. If P and $P \setminus \{p\}$ do not have the same convex hull, we say that p is an *extreme point*. We denote by $P' \subseteq P_i$ the extreme points in P_i, where P_i denotes the inner points in P. Note that a point might lie on the boundary of the convex hull of a point set without being an extreme point. We say that P is *special* if $|P'| \leq 2$. Recall that for a given covering of a point set Q with non-crossing segments, we always assume that the endpoints of the segments are in Q.

Lemma 4.10. *Let P be a set of n points that is not special. Given a covering K of P_i with s non-crossing segments, one can compute in*

$\mathcal{O}(n^2)$-time a convex partition Σ of P with at most $4s + 2|P'|$ faces. Moreover every segment in K is the union of some edges in Σ.

Proof. Let $Q \subseteq P_i$ be the set of the endpoints of segments in the covering. Note that $|Q|$ is at most $2s$. As P is not special, there exist triangulations of Q. We compute a constrained triangulation (for example Delaunay) of Q with respect to the segments of the covering. This can be done in $\mathcal{O}(n \log n)$-time [18], and there are at most $2|Q|$ faces. We observe that the triangulation of Q gives a convex partition of P_i. We add all segments between consecutive points on the boundary of the convex hull of P. Now, it remains to deal with the surface that is within the convex hull of P, but not within the convex hull of P_i. To do that, we add for each point in P' at most two edges to points on the boundary of the convex hull of P. We do it such that the angle between any consecutive edges around a point in P' is at most π. This takes $\mathcal{O}(n^2)$ time [52]. We have now obtained a convex partition of P. □

If one is interested in a convex tiling instead of a convex partition in Lemma 4.10, note that it is possible to add only one edge for each point in P', resulting in a convex tiling with at most $4s + |P'|$ faces.

Lemma 4.11. *Let P be a set of n points. Given a convex tiling Σ of P with f faces, one can compute in $\mathcal{O}(n)$-time a covering K of P_i with at most $6f - 2|P'|$ non-crossing segments. Moreover every segment in K is the union of some edges in Σ.*

Proof. The proof is illustrated in Figure 4.4. Let us denote by $G_0 = (V_0, E_0)$ the plane graph corresponding to the convex tiling, where a point in V_0 is extreme or has degree at least 3. Observe that some points in V_0 might not be in P. Also, the relative interior of an edge in E_0 might overlap with points in P. We assume that G_0 is given with a doubly connected edge list (DCEL) structure. If there is an edge between two points on the boundary of the convex

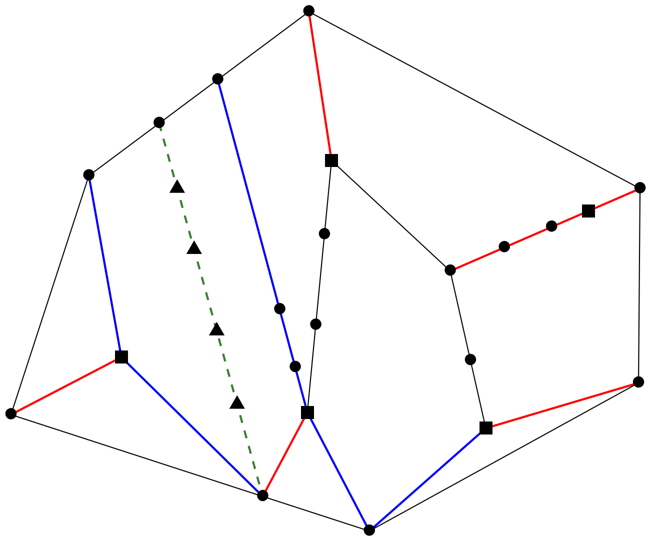

Figure 4.4: Illustration of Lemma 4.11. The green dashed edge and the triangle points are removed at the beginning for the analysis, and added back at the end. The extreme points in P' are represented as square points. The edges in E' are in red. The other edges from P' to the boundary of the convex hull are in blue.

hull of V_0, but not consecutive, we remove this edge. Note that this decreases the number of faces by 1, and does not break the convexity property. We denote by m the number of such edges that we have removed. We also remove from P all points contained in the relative interior of an edge between two points on the boundary of the convex hull. We denote by P'' the extreme points in P_i that we have not removed. As an edge contains at most two points in P', we have $|P''| \geq |P'| - 2m$. Using the DCEL structure, this can be done in $\mathcal{O}(n)$-time. We have obtained a new graph $G = (V, E)$, and there are $f - m$ convex faces in G. We denote by Q the set of inner points that are of degree at least 3 in G. We set $k := |Q|$. Now observe that for each point p in P'', there exists at least one edge e in E with one endpoint in Q, one endpoint on the boundary of the convex hull, such that e overlaps with a point in P''. This is because if we consider p and the two lines going through p and one of the two consecutive vertices in $'P'$ (the one before p and the one after p when going around P'' in clockwise order), they define a wedge in which one edge must lie because of convexity. The point in P'' can be an endpoint of e or in its relative interior. If for a point $p \in P''$ there are several edges that satisfy the conditions, we choose one arbitrarily. We denote these edges by E'. An edge in E' overlaps with exactly one point in P'', thus $|E'| = |P''|$. We denote by E_b the edges not in E' that have a point on the boundary of the convex hull and the other in Q, and we denote $|E_b|$ by m'. The vertices on the boundary of the convex hull are adjacent to two other vertices on the boundary of the convex hull. Moreover, those vertices are incident to $|P''| + m'$ additional edges. We have $2|E| = \sum_{v \in V} \deg(v) \geq 3k + 2(n-k) + |P''| + m' = k + 2n + |P''| + m'$. By Euler's formula, we have $f - m = |E| - n + 1 \geq \frac{k+|P''|+m'}{2} + 1$.

Now, the solution consists of the union of all edges in E incident to two points in Q, with the m edges in E_0 that we have removed, and with the $|P''|+m'$ edges in $E' \cup E_b$. We may need those edges as they might overlap with points in P_i. Note that there are at most $3k$ edges in E incident to two points in Q as G is plane. Moreover, all points in P_i are indeed covered by the edges in our solution. Thus, we obtain a covering of P_i with s segments, where $s \leq 3k + m + m' + |P''| \leq 3(2(f-m) - |P''| - m') + m + m' + |P''| \leq 6f - 5m - 2|P''| \leq 6f - 5m - 2(|P'| - 2m) \leq 6f - 2|P'|$. □

It is now possible to combine Lemmas 4.10 and 4.11 to prove Theorem 4.6.

Proof of Theorem 4.6. Let us denote by f'_m the minimum number of convex faces in a convex tiling of P. We have $f'_m \leq f_m$. First, if P is special, then the three problems are trivial to solve. It remains to prove statement 1. Recall that we assume that P_i is not empty. As P is special, we need exactly one segment to cover the inner points, and thus $s_m = 1$. Now observe that we need between two and four convex faces in a convex partition, and exactly two convex faces in a convex tiling, thus it holds $\frac{s_m}{6} \leq f'_m \leq f_m \leq 8s_m$.

Let us now assume that P is not special. Starting with a covering of the set P_i with s_m non-crossing segments, Lemma 4.10 indicates that it is possible to find a convex partition with at most $4s_m + 2|P'|$ convex faces. Therefore we have $4s_m + 2|P'| \geq f_m$. Starting from a convex tiling with f'_m convex faces, we know from Lemma 4.11 that there exists a covering of P_i with at most $6f'_m - 2|P'|$ non-crossing segments. This implies $6f'_m - 2|P'| \geq s_m$. Note that any segment in a covering can cover at most two points in P'. Therefore we have $s_m \geq |P'|/2$. Putting everything together, we obtain $\frac{s_m}{6} \leq f'_m \leq f_m \leq 8s_m$.

Let us consider a covering of P_i with $s \leq \lambda s_m$ non-crossing segments. By Lemma 4.10, we can compute in $\mathcal{O}(n^2)$-time a convex partition of P with at most $f := 4s + 2|P'|$ faces. We now have $f = 4s + 2|P'| \leq 4\lambda s_m + 2|P'| \leq 4\lambda(6f'_m - 2|P'|) + 2|P'| \leq 24\lambda f'_m$. This implies that the convex partition we have is a 24λ-approximation for MCP and for MCT.

Let us consider a convex tiling of P_i with $f \leq \lambda f_m$ convex faces. Note that this encompasses the case where the convex tiling is actually a convex partition. By Lemma 4.11, we can compute a covering of P_i with at most $s := 6f - 2|P'|$ segments. We have $s = 6f - 2|P'| \leq 6\lambda f_m - 2|P'| \leq 6\lambda(4s_m + 2|P'|) - 2|P'| \leq 24\lambda s_m + 10\lambda|P'| \leq 44\lambda s_m$. □

4.2.2 Approximation algorithms for CPNCS

In this section we present several approximation algorithms for CPNCS. Let us first consider the ones whose approximation ratio is not output-dependent. The best algorithms in terms of approximation ratio are constant-approximation algorithms. The fastest algorithms take quadratic time. Therefore by 2. of Theorem 4.6, all the algorithms we present for CPNCS can be used to obtain approximation algorithms for MCP and MCT with the same order of approximation ratio, and the same order of running time. We have also one algorithm for CPNCS which realises an $\mathcal{O}(\log OPT)$-approximation in time $\mathcal{O}(n^8)$, where OPT denotes the minimum number of segments needed. Using 1. and 2. of Theorem 4.6, we also derive from it the $\mathcal{O}(\log OPT)$-approximation algorithm for MCP and MCT running in time $\mathcal{O}(n^8)$, where now OPT denotes the minimum number of faces needed in a convex partition, or in a convex tiling, respectively. This is how we prove Theorem 4.5. We first present an easy approximation algorithm running relatively fast, at the cost of a high approximation ratio.

Theorem 4.12. *There exists an $\sqrt{n}\log(n)$-approximation algorithm for CPNCS running in $\mathcal{O}(n^2)$-time.*

Proof. Let us denote by ℓ_m the minimum number of lines needed to cover P, as in the CPL problem. We denote by s_m the minimum number of segments in a valid solution of CPNCS. We have $\ell_m \leq s_m$.

Using the greedy algorithm for set cover problems, we can compute a covering of P with ℓ lines, where $\ell \leq \log(n)\ell_m$ [47, 59]. The greedy algorithm runs in quadratic time. Indeed, it is folklore that the greedy algorithm for covering a set X with the family of subsets $\mathcal{F} \subseteq 2^X$ can be implemented in $\mathcal{O}(Z)$ time, where $Z = \sum_{F \in \mathcal{F}} |F|$. In our situation, Z is the number of point-line incidences, and so $Z = \mathcal{O}(n^2)$ because each point lies on at most $n-1$ lines.

We distinguish two cases, depending on the value of ℓ. We denote by $\phi(n)$ a threshold function, that will be determined later. In the first case, we assume $\ell \geq \phi(n)$. In this situation, we cover each point in P by a segment reduced to that point. We have n segments, and we needed at least $s_m \geq \ell_m \geq \ell / \log(n) \geq \phi(n)/\log(n)$. The approximation ratio is $\frac{n \log(n)}{\phi(n)}$.

Now, let us assume $\ell < \phi(n)$. We transform each of the ℓ lines into a segment, such that the new segments still cover P. Now, at each of the $\mathcal{O}(n^2)$ intersections between the relative interior of a pair of segments, we split one segment into two, such that there is no crossing anymore. Let us denote by s the number of segments obtained. We have $s \leq \ell^2 \leq \ell \log(n)\ell_m \leq \phi(n)\log(n)s_m$. The approximation ratio is $\phi(n)\log(n)$.

We make the two approximation ratios equal by setting $\phi(n) := \sqrt{n}$. We obtain a $\sqrt{n}\log(n)$-approximation. □

Mitchell presented in a technical report some approximation algorithms for the problem of covering a point set with a minimum number of pairwise-disjoint triangles [63]. In his problem, the triangles of the covering must be subtriangles of some triangles given as input, for otherwise the problem would be trivial. He makes the assumption that no three points are on a line. We adapt his algorithms to our setting of CPNCS for point sets with no constraint. It seems that there were two mistakes in his proof, that we show how to fix.

Let P be a set of n points. By doing a rotation if necessary, we can assume that no two points in P have the same x-coordinate. We say that a trapezoid is *constrained* if 1) it has two disjoint vertical sides, each lying on a line that contains a point in P, and 2) the two remaining sides are lying on lines that contain each at least two points in P. Note that there are $\mathcal{O}(n^6)$ constrained trapezoids.

In his paper, Mitchell calls the trapezoids "canonical" instead of "constrained" [63]. We make the choice of changing the name for better clarity later. Also, concerning constraint 1), he has the stronger constraint that the vertical sides must each contain a point in P. It seems to be a mistake, for otherwise it is not clear how his dynamic programming algorithms work, and some of his arguments do not hold. Anyway, even with his definition, he only uses the fact that there are $\mathcal{O}(n^6)$ constrained trapezoids for computing the running time of his algorithms. Therefore there is no loss in using our definition.

We also allow for some degeneracies. Let us consider a triangle with vertices a, b and c, not all three on a line. If a is in P, the segment with endpoints b, c is vertical and lies on a line that contains a point in P, and the segments with endpoints a, b and a, c respectively are contained in some lines ℓ and ℓ' such that ℓ and ℓ' contains at least two points in P, then we say that the triangle is a constrained trapezoid. If a constrained trapezoid is split into two halves by a vertical line ℓ going through its interior, with ℓ containing a point in P, we obtain two constrained trapezoids. Likewise, if a segment s is in a constrained trapezoid τ, such that s lies on a line that contains at least two points in P, s intersects the interior of τ, and the endpoints of s are contained in the vertical sides of τ, then s splits τ into two constrained trapezoids.

For a set of points P where no two points have the same x-coordinate, we define the *enclosing trapezoid* as follows. Let ℓ_1 be the vertical line that contains the leftmost point in P, and let ℓ_2 be the vertical line that contains the rightmost point in P. Let L be the set of all lines containing at least two points in P. Observe that no line in L is vertical. We denote by a the highest intersection point between ℓ_1 and a line in L. We denote by b the lowest point intersection point between ℓ_1 and a line in L. Similarly, we denote by c and d, respectively, the highest intersection point, respectively the lowest intersection point, between ℓ_2 and a line in L. We denote by ℓ_3 the line containing a and c, and by ℓ_4 the line containing b and d. The *enclosing trapezoid* of P is the constrained trapezoid of $P \cup \{a, b, c, d\}$ defined by ℓ_1, ℓ_2, ℓ_3 and ℓ_4. It is denoted by $\mathcal{T}_P$.

Mitchell uses in his paper the notion of *guillotine property*. He shows that if there is a covering of the points in P with s elements, then there is a covering of P with at most $\mathcal{O}(s \log s)$ elements having the guillotine property. He then presents an algorithm, and claims that it outputs an optimal solution among all coverings that have the guillotine property. While we agree that his algorithm outputs a solution with at most $\mathcal{O}(s \log s)$ elements, we present a counterexample to the fact that his algorithm outputs an optimal solution among all coverings that have the guillotine property. Although he considers the problem of covering points with triangle, he reduces the problem to covering a set of points with constrained trapezoids. He defines the guillotine property for trapezoids as follows: A set $\mathcal{T}$ of constrained trapezoids has the guillotine property if a) it contains at most one trapezoid, or if b) there exists a partitioning line ℓ containing at least two points in P not intersecting the interior of any constrained trapezoid in $\mathcal{T}$, such that the sets of constrained trapezoids on both sides of ℓ also have the guillotine property, or if c) there exists a vertical partitioning line ℓ not intersecting the interior of any constrained trapezoid in $\mathcal{T}$, such that the sets of constrained trapezoids on both sides of ℓ are not empty also have the guillotine property [63]. Mitchell's wording is not exactly the same as ours but the two definitions are equivalent.

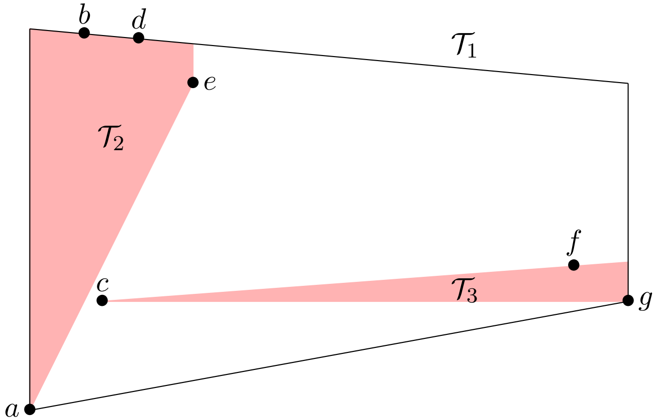

Figure 4.5: The constrained trapezoids $\mathcal{T}_2$ and $\mathcal{T}_3$ have the guillotine property, and cover the points in P. Mitchell's algorithms applied to $\mathcal{T}_1$ outputs a solution with three trapezoids and is therefore not optimal.

In Figure 4.5 we have represented three constrained trapezoids, with $\mathcal{T}_3$ being reduced to a triangle. Observe that $\mathcal{T}_2$ and $\mathcal{T}_3$ have the guillotine property. Indeed the line going through a and c satisfies condition b) of the guillotine property. Let us now apply Mitchell's algorithm to the constrained trapezoid $\mathcal{T}_1$. In Mitchell's setting, not all constrained trapezoids can be used to cover the points in P: they must be subtrapezoids of some given trapezoids. Here the two given trapezoids are $\mathcal{T}_2$ and $\mathcal{T}_3$. Observe that a minimum covering of P uses $\mathcal{T}_2$ and $\mathcal{T}_3$. We claim that Mitchell's algorithm outputs at least three trapezoids. His algorithm recurses on all the ways of splitting P with a vertical line. Observe that any vertical line going through a point in P intersects the interior of $\mathcal{T}_2$ or $\mathcal{T}_3$, and therefore cannot lead to a solution with two trapezoids. In addition, his algorithm recurses on all the ways of splitting P with a segment that contains at least two points in P, and whose endpoints are on the vertical sides of $\mathcal{T}_1$. Observe that any such segment σ splits the interior of $\mathcal{T}_2$ or $\mathcal{T}_3$. Moreover, the points contained in the trapezoid that is split by σ are not on the same side of σ. This implies that all those recursions will lead to solutions with at least three trapezoids. Figure 4.5 thus depicts a counterexample to the fact that Mitchell's algorithm outputs an optimal solution among the ones that have the guillotine property. The reason for that is that the line ℓ going through a and e, and the line ℓ' going through a and c, which are the certificates that $\mathcal{T}_2$ and $\mathcal{T}_3$ have the guillotine property, do not intersect $\mathcal{T}_1$ only at its vertical sides. Therefore splitting along ℓ and ℓ' is not tested by Mitchell's algorithm.

We define the *strong guillotine property* in the special case of segments. We show that if there is a covering of P with s non-crossing segments, then there is a covering of S with $\mathcal{O}(s \log s)$ non-crossing segments having the strong guillotine property. We then present an algorithm that outputs an optimal solution among all the coverings with non-crossing segments having the strong guillotine property. Let S be a set of non-crossing segments covering P. We assume that the endpoints of the segments in S are in P. We say that S has the strong guillotine property with respect to a constrained trapezoid $\mathcal{T}$ that contains all segments in S if a) S contains at most one segment, or if b) there exists a partitioning line ℓ containing at least two points in P and at least one segment in S, such that for any segment $s \in S$, ℓ either contains s or does not intersect the relative interior of s, and ℓ splits $\mathcal{T}$ into two constrained trapezoids $\mathcal{T}_1$ and $\mathcal{T}_2$, such that the segments in $\mathcal{T}_1$, respectively $\mathcal{T}_2$, have the strong guillotine property with respect to $\mathcal{T}_1$, respectively $\mathcal{T}_2$, or if c) there exists a vertical line not intersecting with the relative interior of any segment in S, that splits $\mathcal{T}$ into two constrained trapezoids $\mathcal{T}_1$ and $\mathcal{T}_2$, such that the segments in $\mathcal{T}_1$, respectively $\mathcal{T}_2$, have the strong guillotine property with respect to $\mathcal{T}_1$, respectively $\mathcal{T}_2$. Observe that the line ℓ in case b) only intersects the vertical sides of $\mathcal{T}$, for otherwise ℓ would not split $\mathcal{T}$ into constrained trapezoids. We simply say that

S has the strong guillotine property if it has the strong guillotine property with respect to the enclosing trapezoid $\mathcal{T}_P$.

Lemma 4.13. *If there exists a covering of P with s non-crossing segments, then there exists a covering of P with $\mathcal{O}(s \log(s))$ non-crossing segments with the strong guillotine property.*

Proof. Recall that we assume that the endpoints of the segments are in P, by cropping them if need be. We can even crop some segments further such that they are pairwise-disjoint (it may be that now some segments are reduced to points). Consider the endpoints of the segments in that covering, that we denote by P'. We denote $|P'|$ by n', and we have $n' \leq 2s$. Note that no two points in P' have the same x-coordinate. We denote by X the set of x-coordinates of the points in P'. We now consider the segment tree based on X, as defined in [69]. The segment tree defines some canonical intervals. Each interval, whose endpoints are in X, is partitioned into $\mathcal{O}(\log s)$ canonical intervals. We partition each segment in the covering, such that the projection on the x-axis of each new segment is a canonical interval. Therefore we obtain a covering of P with $\mathcal{O}(s \log(s))$ non-crossing segments. We claim that this family of segments has the strong guillotine property. Let us denote by x_i, $1 \leq i \leq n'$ the elements in X, ordered by increasing value. We distinguish two cases. If there exists a segment σ whose projection on the x-axis is equal to the interval $[x_1, x_{n'}]$, then we recurse on the parts above and below σ which contain some segments. Observe that if $n' = 2$ we are done. If there is no such segment, then by definition of a segment tree, there is no segment in the covering whose relative interior intersects the vertical line ℓ with x-coordinate equal to $x_{\lfloor (1+n')/2 \rfloor}$. Thus we can recurse on the left and right side of ℓ. □

Theorem 4.14. *There exists an $\mathcal{O}(\log(OPT))$-approximation algorithm running in $\mathcal{O}(n^8)$-time for CPNCS.*

Proof. We explain how to recursively compute a minimum covering of P with non-crossing segments under the constraint that the solution has the strong guillotine property. The approximation ratio for the CPNCS problem when this additional constraint is removed follows from Lemma 4.13. If P is empty, we return no segment, which is a valid solution. If P can be covered with a single segment, we return that segment. This can be tested in $\mathcal{O}(n^2)$ time using duality. Now let us assume that not all points in P are on a line. We compute the enclosing trapezoid $\mathcal{T}_P$ of P. We consider the four vertices a, b, c, d of $\mathcal{T}_\mathcal{P}$. We start by adding the segment with endpoints a, c, and the segment with endpoints b, d. Now all the points to cover are within the enclosing trapezoid $\mathcal{T}_P$. We distinguish two cases, according to whether a segment with endpoints on the vertical sides of $\mathcal{T}_P$ is in a minimum covering with non-crossing segments having the strong guillotine property. If it is, we can add it to the solution and recurse on the two new constrained trapezoids. If no such segment is part of a minimum solution, then there exists a vertical line ℓ that splits a minimum solution into two parts, such that ℓ does not intersect the relative interior of any segment in that minimum solution. We can recurse on the $\mathcal{O}(n)$ choices of splitting vertically the constrained trapezoid into two constrained trapezoids. For each of the $\mathcal{O}(n^2)$ recursions, we compute the number of segments corresponding to that solution, and we output the solution corresponding to the one that minimises the number of segments.

To optimise we can do dynamic programming, and solve first the thinnest constrained trapezoids (in terms of width on the x-axis). There are $\mathcal{O}(n^6)$ constrained trapezoids, and we take quadratic time for each of them, so the total running time is $\mathcal{O}(n^8)$. □

It is possible to use the segment tree technique for the computation as done by Mitchell. It reduces the running time to $\mathcal{O}(n^7)$ at the cost of a slightly worse approximation ratio.

Theorem 4.15. *There exists an* $\mathcal{O}(\log(n))$*-approximation algorithm running in* $\mathcal{O}(n^7)$*-time for CPNCS.*

Proof. We consider the set X of x-coordinates of points in P. We compute the corresponding segment tree in $\mathcal{O}(n \log(n))$-time. Let us consider a minimum covering of P with s non-crossing segments. We crop the segments so that their endpoints are in P. We partition each segment in the covering such that the projection of each new segment on the x-axis is a canonical interval of the segment tree. We say that a segment is *canonical* if its projection on the x-axis is a canonical interval. We observe that there is a covering of P with $\mathcal{O}(s \log(n))$ non-crossing canonical segments. Thus, we can adapt the algorithm of Theorem 4.14 to obtain an $\mathcal{O}(\log(n))$-approximation, by outputing a minimum covering of P with non-crossing canonical segments.

We call a constrained trapezoid whose projection on the x-axis is a canonical interval a *canonical trapezoid* (note that this definition is not the same as Mitchell's). As there are $\mathcal{O}(n)$ canonical intervals, there are $\mathcal{O}(n^5)$ canonical trapezoids. We do as in the algorithm of Theorem 4.14. The difference is that when we assume that there exists a vertical line ℓ that splits P, with ℓ not intersecting the relative interior of any segment in an optimal solution, we can assume that the x-coordinate of ℓ is equal to the median of X. For each canonical trapezoid we still do $\mathcal{O}(n^2)$ recursions, so the overall running time of the algorithm is $\mathcal{O}(n^7)$. □

We say that a point set P is k*-directed* if there exists a set D of k directions, such that for any line ℓ that contains at least three points in P, the direction of ℓ is in D. Without loss of generality, we can assume that the two directions of a 2-directed point set P are vertical and horizontal. Indeed, this has no impact for the CPNCS problem. For convenience, for any set of directions D and any segment s reduced to a point, we say that the direction of s is in D. We say that a set of segments S has the *autopartition property* if $|S| = 1$, or if there exists a line ℓ which contains at least one segment in S, and splits S into two sets that are either empty or have the autopartition property. The relative interior of a segment in S is either contained in ℓ or does not intersect ℓ.

Lemma 4.16. *Let P be a k-directed point set with set of directions D. If there exists a covering of P with s non-crossing segments, then there exists a covering of P with $\mathcal{O}(sk)$ non-crossing segments having the autopartition property, such that the direction of each segment is in D. If $k = 2$ then there exists a covering with at most $4s$ non-crossing segments, being vertical or horizontal, having the autopartition property.*

Proof. Let D be the set of k directions, such that the direction of any line that contains at least three points in P belongs to D. From the covering with s segments, we can obtain a covering with at most $2s$ segments such that the direction of each segment is in D. Indeed, a segment in the covering whose direction is not in D contains at most two points in P. We crop some segments if necessary such that no two segments intersect, and they still cover P. Now, we use a theorem by Tóth who showed that any set of s' disjoint segments having up to k directions have an autopartition of size $\mathcal{O}(s'k)$ [78]. This immediately implies the result.

Let us now assume $k = 2$. Let us consider a set of s' segments that are vertical or horizontal. There exists a partition of the segments that contains at most $2s'$ segments, and which has the autopartition property [26]. An upper bound of $3s'$ was first shown by Paterson and Yao [66]. It was then improved to $2s'$ by d'Amore and Franciosa [21], although not explicitly. Dumitrescu, Mitchell and Sharir made the result explicit later [26]. □

Theorem 4.17. *There exists an $\mathcal{O}(k)$-approximation algorithm for CPNCS in k-directed sets running in $n^{\mathcal{O}(k)}$. Furthermore, there ex-*

ists a 4-approximation algorithm for CPNCS in 2-directed sets running in time $\mathcal{O}(n^5)$.

Proof. Let P be a k-directed point set with set of directions D. We show how to compute an optimal covering of P with non-crossing segments, such that the solution has the autopartition property, and the direction of each segment in the solution is in D. By Lemma 4.16, our algorithm realises an $\mathcal{O}(k)$-approximation, and a 4-approximation for the special case $k = 2$.

The recursion of the algorithm is as follows. If there exists a segment that contains all points in P, we add it to the solution. This can be tested in time $\mathcal{O}(n^2)$. Otherwise, for each direction δ in D, we recurse in the $\mathcal{O}(n)$ ways of splitting P with a line ℓ, such that the direction of ℓ is δ, and ℓ contains at least a point in P. We add to the solution the shortest segment containing the points in ℓ, and we recurse on both sides of ℓ.

The subsets of $\mathbb{R}^2$ we are considering in the recursion are defined by giving for each direction the two extreme points for that direction. For each point we have $\mathcal{O}(n)$ choices, so in total there are n^{2k} of such subsets. For each subset by using dynamic programming, we need $\mathcal{O}(nk)$-time. Thus, the total running time is $n^{\mathcal{O}(k)}$, and simply $\mathcal{O}(n^5)$ for the special case $k = 2$. □

A natural question is whether we can use the autopartition property when the number of directions is not fixed. It is known that a set of s pairwise-disjoint segments allows for an autopartition with $\mathcal{O}(s \log s / \log\log s)$ segments [79]. This is tight [77]. Using the techniques we have presented, one could hope to obtain an $\mathcal{O}(\log OPT / \log\log OPT)$-approximation algorithm. However, this autopartition might not have any good structure, and so we cannot use dynamic programming because there are too many subsets of $\mathbb{R}^2$ to consider. In any case, because of the tightness on the number of segments in the autopartition, it seems that the autopartition technique cannot be used to obtain constant-approximation algorithms, or even an $o(\log OPT/ \log\log OPT)$-approximation algorithm.

4.3 Fixed-parameter algorithm for CPNCS

As mentioned in the introduction, there are known fixed-parameter algorithms for MCP, where the parameter is the number of inner points. We present here a fixed-parameter constant-approximation algorithm for MCP and MCT, where the parameter is the number of faces in a minimum convex partition or a minimum convex tiling, respectively. For point sets where no three points are on a line, the minimum number of convex faces is at least half the number of inner points [52], and the number of convex tiles is at list a sixth of the number of inner points [25]. However, as shown in Figure 4.2, when we allow for several points to be on a line, the number of inner points can be arbitrarily larger than the number of convex faces in a minimum convex partition. If the number of inner points is significantly higher than the number of convex faces needed, our algorithm has a lower running time. We first show that CPNCS is in FPT.

Theorem 4.18. *It is possible to compute in $\mathcal{O}(2^{k^2}k^{7k} + n^4 \log n)$-time whether a point set P can be covered with at most k non-crossing segments, and to output such a covering if it exists.*

The proof uses a kernelisation technique presented by Langerman and Morin for CPL [55]. Assume there is a line ℓ that contains at least $k+1$ points in P. Then in any covering of P with at most k lines, ℓ must be in the covering. Otherwise, we would need at least $k+1$ lines to cover the points contained in ℓ. Now one can compute all of these lines that contain at least $k+1$ points, dismiss all of the covered points, until no line covers more than k of the remaining points. If there remains more than k^2 points, then there is no covering of the point set with at most k lines. Otherwise, one can compute every way of covering the $\mathcal{O}(k^2)$ remaining points, and check whether there is one that uses in total at most k lines. In our setting, we are looking for a covering with non-crossing segments, which makes it more difficult. Indeed, if a line ℓ contains at least $k+1$ points, we only know that ℓ must contain at least one segment of the covering. This means that we cannot simply dismiss the points covered by such a line. Also, we have to be careful about crossings. Before proving Theorem 4.18, we first show several lemmas. For a point set P, we say that a segment s is a *P-segment* if its endpoints are in P. Recall that we only consider coverings of a point set P with non-crossing P-segments.

Definition 4.19. Let P be a point set, and let s and t be two crossing P-segments. We denote by p the intersection of s and t. We determine four points in P, that we call the *points enclosing p*. There are two points on $s \cap P$ and two points on $t \cap P$. The two points on $s \cap P$, denoted by u and v, are such that the segment with endpoints u and v, which we denote by uv, is the shortest P-segment contained in s whose relative interior contains p. Likewise, the two points u' and v' are such that $u'v'$ is the shortest P-segment contained in t whose relative interior contains p. The points u, v, u' and v' are the points enclosing p.

Lemma 4.20. *Given a set P of n points, it is possible to compute in time $\mathcal{O}(n^4 \log n)$ the pairs of crossing P-segments, to find whether their intersection p is in P, and to store the points enclosing p. Additionally, we can also store for each P-segment how many points in P they contain, and the list of those points.*

Proof. Let s be a P-segment. We first store the number of points contained in s, as well as the list consisting of those points. We then sort the list so that when going from one endpoint of s to the other, the points appear consecutively on the list. As there are $\mathcal{O}(n^2)$ of such segments, this preprocessing can be done in time $\mathcal{O}(n^3 \log n)$. Let s and s' be some crossing P-segments. There are $\mathcal{O}(n^4)$ pairs of such segments. We denote by p the intersection of s and s'. We check whether p is in P. This can be done in time $\mathcal{O}(\log n)$ by searching through the list of points in s. We denote by u, v, u' and v' the points enclosing p. Observe that given s and s', it takes time $\mathcal{O}(\log n)$ to find the points enclosing p, and to test whether p is in P. Thus when considering all pairs of segments, we can compute this information in time $\mathcal{O}(n^4 \log n)$, and so given the endpoints of some crossing P-segments s and s', we can retrieve this information in constant time. Thus, the total running time of the algorithm is in $\mathcal{O}(n^4 \log n)$. □

Lemma 4.21. *Given a set P of n points, and a natural number k, it is possible to find in time $\mathcal{O}(2^{k^2} + n^4 \log n)$ either a certificate that there is no covering of P with at most k non-crossing segments, or to output a family $\mathcal{F}$ of $\mathcal{O}(2^{k^2})$ sets S containing at most k non-crossing P-segments, with the following properties: For any fixed covering of P with at most k non-crossing P-segments, there exists a set S in $\mathcal{F}$ such that a) any segment $s \in S$ contains at least $k+1$ points in P, b) for each segment t of the covering, if $|P \cap t \cap s| \geq 2$ for some $s \in S$, then t is contained in s, and c) if a segment of the covering contains at least $k+1$ points in P, then it is contained in a segment in S.*

Let P be a point set and let k be a natural number. Observe that if a set S of segments satisfies property a), then in a covering with at most k segments of P, each segment s in S contains at least one segment t of the covering, such that $|P \cap t| \geq 2$. Indeed if there exists a segment $s \in S$ such that for any segment t in the covering, we have that $s \cap t$ contains at most one point in P, then at least $k+1$ segments are needed to cover the points in $P \cap s$. This implies that if S consists of m segments and satisfies properties a) and b), then there are at least m segments in the considered covering of P

with non-crossing segments.

Proof of Lemma 4.21. We first do some preprocessing by using the algorithm of Lemma 4.20. This takes $\mathcal{O}(n^4 \log n)$ time. We create a list L of segments, which at the beginning is empty, and will contain the segments in S when we are done. For each line ℓ that contains at least $k+1$ points, we find the extremal points p and q of P contained in ℓ in time $\mathcal{O}(n)$. Then we add the line segment with endpoints p and q to L. Using the algorithm presented by Guibas *et al.* [38], we can compute all lines containing more than k points in time $\mathcal{O}(\frac{n^2}{k} \log(\frac{n}{k}))$. If there are more than k of such lines, we already know that there is no covering of P with at most k non-crossing segments of P. Indeed such a covering can only exist if there exists a covering of P with at most k lines. Let us now assume that there are at most k such lines. We add all corresponding segments to L in total time $\mathcal{O}(kn + \frac{n^2}{k} \log(\frac{n}{k}))$. Let us show that the segments in L satisfy properties a), b) and c), although they might still be crossing. First, property a) holds by definition. Moreover property b) holds for all covering of P with at most k segments because a segment in L containing points p and q also contains all points on the line (p, q). Finally, property c) also holds trivially for all covering of P with at most k segments.

We are now going to modify L and make copies of it while maintaining the fact that properties a), b) and c) hold. Our aim is that no two segments in L cross. Let us consider one segment s in L which is crossed by another segment s' in L. We denote by p the intersection of s and s'. We retrieve the points u and v such that uv is the shortest P-segment in s whose relative interior contains p. We do likewise with u' and v' in s'. Observe that not both uv and $u'v'$ can be in a covering of P with non-crossing segments. More generally, in a valid covering, at least one of uv and $u'v'$ is not contained in any segment of the covering. We create one copy of L, and recurse on two cases, one where we assume that uv is not contained in a segment of the covering, and one where we assume that $u'v'$ is not contained in a segment of the covering. In the case where we assume that uv is not in the covering, s' remains in L, and might still be removed at a later step. Let us assume for now that uv is not contained in a segment of the covering. We remove s from L. The candidate segment s' splits s at p into two sides. Let us denote by x and y the endpoints of s, with u being closer to x than v is. If p is not in P, we consider the segments xu and vy. If p is in P, we consider the segments xp and py. Any of the two new segments that contains more than k points in P is added to L. Indeed property a) holds by definition. Moreover property b) holds because s was in L, and we are assuming that the segment uv is not contained in a segment of the covering. If a segment contains at most k points, we do not add it to L. We claim that property c) still holds. This is because if a point $q \in P$ which lies on a line that contains more than k points is not contained in some segment in L, that means that if a segment t contains q as well as at least k other points in P, then t also contains some segment which we are assuming not to be contained in a covering.

If we obtain more than k segments in L, we stop this branch of the recursion, as we already know that there is no valid covering of P with at most k segments, assuming that uv is not contained in a segment of the covering. We now iterate over all crossing segments in L. We obtain $O(k)$ segments in L, which are by construction non-crossing. As the depth of the recursion tree is in $\mathcal{O}(k^2)$, the number of leaves is in $\mathcal{O}(2^{k^2})$. We would like to say that each recursion implies the existence of one more segment in a covering with non-crossing segments, but this is a priori not the case. Therefore, if the number of lines containing more than k points is in $\Omega(k)$, we might have to do $\Omega(k^2)$ recursions. We can do the computation in total time $\mathcal{O}(2^{k^2} + kn + \frac{n^2}{k}\log(\frac{n}{k}))$, using the information we preprocessed. If we add to it the running time of the preprocessing, the total running time of the algorithm is in $\mathcal{O}(2^{k^2} + n^4 \log n)$. $\square$

Proof of Theorem 4.18. We first use the algorithm of Lemma 4.21. In particular, we keep the information that is preprocessed with the algorithm of Lemma 4.20. If we have a certificate that there is no covering of P with at most k non-crossing segments, we stop. Let us assume that the algorithm outputs a family $\mathcal{F}$ of $\mathcal{O}(2^{k^2})$ sets S, such that $\mathcal{F}$ satisfies the conditions of Lemma 4.20. Let us consider a fix covering of P with at most k non-crossing segments, assuming one exists. We guess the corresponding set of segments S in time $\mathcal{O}(2^{k^2})$. We call the segments in S *candidate segments.*

Let us denote by m the number of candidate segments, and by m' the number of points not contained in a candidate segment. Computing m and m' takes $\mathcal{O}(n)$ time. If m' is larger than k^2, we output that there is no solution. Indeed, by property c), a segment in the covering can contain at most k points which are not contained in some candidate segment. If $m + m'$ is at most k, then we take the covering consisting of all candidate segments, and segments reduced to a point for each point which is currently uncovered. Otherwise, for a covering to have at most k non-crossing segments, there must be a segment that contains at least two points which are currently not covered by candidate segments. Indeed by properties a) and b) we know that for each candidate segment s, there exists a segment t in the covering that is contained in s, and therefore t contains no point currently uncovered. Thus to have a covering with fewer than $m + m'$ segments, there must be a segment σ in the covering which contains at least two points for now uncovered. Recall that m' is at most k^2. We consider the $\mathcal{O}(k^4)$ lines going through at least two uncovered points. As we have shown, there exists a line ℓ that contains two uncovered points, and also contains σ. Observe that the endpoints of σ might be contained in some candidate segments. We first guess in $\mathcal{O}(k^4)$ time the largest segment τ contained in σ, such that the endpoints of τ are uncovered points. If ℓ intersects a candidate segment at a point $p \in P$, such that p is not in τ, we want to guess whether p is an endpoint of σ. As there are at most

k candidate segments, given τ, we can guess the endpoints of σ in $\mathcal{O}(k^2)$ time. Therefore, we can guess σ in $\mathcal{O}(k^6)$ time. We find in $\mathcal{O}(k)$ time the list of candidate segments σ intersects, and also check that σ does not intersect any segment that we have already taken in the solution during a past iteration. For each candidate segment s that σ intersect, we do as in the algorithm of Lemma 4.20 and we split s into two sides. We also update m and m'. For a specific candidate segment, this can be done in constant time thanks to the preprocessing. We now iterate from the beginning of the paragraph. At each iteration, we are either done, or we have one more segment in our partial covering. Therefore we iterate at most k times. The total running time of this algorithm (not including the preprocessing) is in $\mathcal{O}(k^{7k})$. The total running time of the algorithm is in $\mathcal{O}(2^{k^2}k^{7k} + n^4 \log n)$. □

Theorem 4.22. *It is possible to compute in time $\mathcal{O}(2^{36f^2}f^{42f+1} + n^4 \log n)$ a convex partition of a point set P with at most $24f$ convex faces, where f denotes the minimum number of convex faces required. The same holds when considering convex tilings.*

Proof. We first compute a minimum covering of the inner points in time $\mathcal{O}(2^{s^2}s^{7s+1} + n^4 \log n)$ by applying the algorithm of Theorem 4.18 for $k = 1, 2, \ldots, s$, where s denotes the minimum number of segments required in a covering of the inner points. Then, by 2. of Theorem 4.6, we obtain in $\mathcal{O}(n^2)$-time a convex partition with at most $24f$ convex faces. The same holds with convex tilings for the same arguments. As by 1. of Theorem 4.6, we have $s \leq 6f$, the total running time of the algorithm is as stated. □

There is a strong similarity between CPNCS and Maximum Independent Set in Segment Intersection Graphs (MISSIG). As an example, we show in Section 4.4.5 that CPNCS is NP-hard by doing a reduction from MISSIG. We have shown that CPNCS is in FPT, but Marx has shown that MISSIG is W[1]-hard [60]. We do a sanity check and explain why his hardness reduction does not apply to CPNCS. In his reduction, there are $f(k)$ gadgets, each gadget containing $\mathcal{O}(n^2)$ segments. In each gadget, a constant number of segments has to be taken in an independent set of size k. We could try to mimic our NP-hardness reduction of Section 4.4.5: Take the same set of segments as Marx, and then replace each segment by a set of four collinear points. Taking a segment in the independent set corresponds to covering these four points with one segment. If a segment is not taken in the independent set, then we need two segments to cover the four points. Therefore, one needs $\Omega(n)$ segments in each gadget to cover the points, and not some constant number, which implies that the W[1]-hardness reduction we are trying to do is not valid.

We now discuss why the FPT algorithm and the techniques presented for CPNCS do not give us an FPT algorithm for MCP where the parameter is the number of faces. One can first use Lemma 4.21 to guess some candidate segments. Then it is possible to guess how many vertices of degree at least 3 lie on each of the $\mathcal{O}(k)$ candidate segments. Then we can enumerate all plane graphs on this vertex set, and guess which one corresponds to our convex partition. However, for now we have only guessed on which candidate segment does a vertex lie. It remains to check whether all those points can be placed at points in P, while preserving the fact that the graph is a convex partition. This can be modelled as an integer linear programming problem, but as the number of constraints is linear in n, this does not give an FPT algorithm.

4.4 Hardness

Our proof builds upon gadgets introduced by Lingas [58]. He used them to prove NP-hardness of several decision problems, including *Minimum Convex Partition for Polygons with Holes* and *Minimum Rectangular Partition for Rectilinear Polygons with Holes*. In the second problem, Steiner points are allowed. However, as noted by Keil [50], one can easily adapt Lingas' proof to not use Steiner points. We use a similar idea to prove NP-hardness of the MCP problem. Lingas' proofs for the two problems are similar, and consist in a reduction from the following variation of planar 3-SAT. The instances are a CNF formula F with set of variables X and set of clauses C, and a planar bipartite graph $G = (X \cup C, E)$, such that there is an edge between a variable $x \in X$ and a clause $c \in C$ if and only if x or $\bar{x}$ is a literal of c. Moreover, each clause contains either two literals or three, and if it contains three, the clause must contain at least one positive and one negative literal. Lingas refers to this decision problem as the *Modified Planar 3-SAT* (MPLSAT). Lingas claims that planar satisfiability can easily be reduced to MPLSAT by adding new variables. As planar satisfiability was shown to be NP-complete by Lichtenstein [57], this would imply that MPLSAT is NP-complete too. For the sake of completeness, we remark that it is not clear why adding these new variables would not break the planarity of the graph. This can be solved by considering the following lemma of Lichtenstein:

Lemma 4.23 (Lichtenstein [57])**.** *Planar satisfiability is still NP-complete even when, at every variable node, all the arcs representing positive instances of the variable are incident to one side of the node and all the arcs representing negative instances are incident to the other side.*

This lemma can easily be strengthened to the case of planar 3-SAT as noted by Lichtenstein, and explicitly done by Tippenhauer [76] in his Master's thesis. From here a reduction to MPLSAT becomes indeed straightforward.

Theorem 4.24. *MPLSAT can be reduced in polynomial time to MCP, and to MCT.*

As it is easy to see that MCP is in NP, Theorem 4.24 implies that MCP is NP-complete. The question whether MCT is in NP is still open. Let P be a point set, and let us consider the set $\mathcal{L}$ of all lines going through at least two points in P. Let P' be the set of points at the intersection of at least two lines in $\mathcal{L}$. One might think that there exists a minimum convex tiling, such that all Steiner points belong to P'. We show in Figure 4.6 that this is not the case.

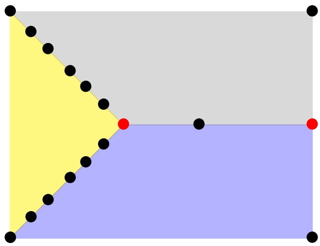

Figure 4.6: Points in black are the input point set P. The points in red are Steiner points in the unique minimum convex tiling of P. The point in red to the right does not belong to the set P'.

We first prove Theorem 4.24 for the MCT problem. We do the reduction from MPLSAT to MCT by constructing a point set in three steps. First we construct a non-simple polygon, in a similar way as in Lingas' proof, with some more constraints. Secondly, we add some line segments to build a grid around the polygon, and finally we discretise each line segment into evenly spaced collinear points. The idea of the first part is to mimic Lingas' proof. The second part makes the correctness proof easier, and the last part transforms the construction into our setting. The aim of the grid is to force the polygons in a minimum convex tiling to be rectangular.

We use the gadgets introduced by Lingas, namely cranked wires and junctions [58]. A wire is shown in Figure 4.7. It consists of a loop delimited by two polygons, one inside the other. In Lingas' construction, the two polygons are simple, and a wire is therefore a polygon with one hole. Moreover in his proof the dimensions of the cranks do not matter. In our case, the polygon inside is not simple, and each line segment has unit length. Each wire is bent several times with an angle of 90°, as shown in Figure 4.7, in order to close the loop.

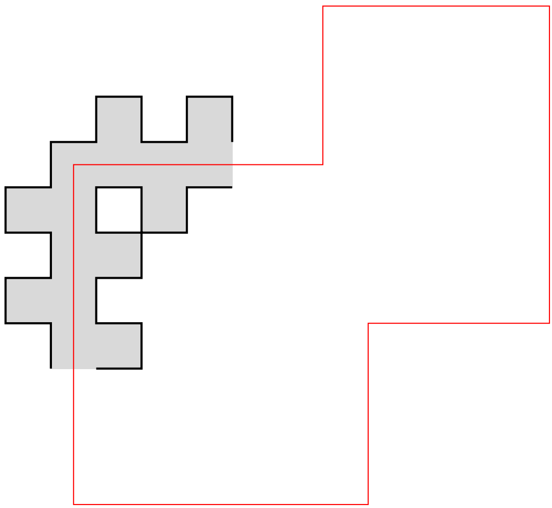

Figure 4.7: A cranked wire, edges are in black and its interior is in grey. The wire follows the whole red loop, but for the sake of simplicity, only a section of the wire at a bend is drawn.

The wires are used to encode the values of the variables, with one wire for each variable. We are interested in two possible tilings of a wire, called vertical and horizontal, which are shown in Figure 4.8.

As in Lingas' proof, we interpret the vertical tiling as setting the variable to *true*, and the horizontal as *false*. Lingas proved the following:

Lemma 4.25 (Lingas [58])**.** *A minimum tiling with convex sets of*

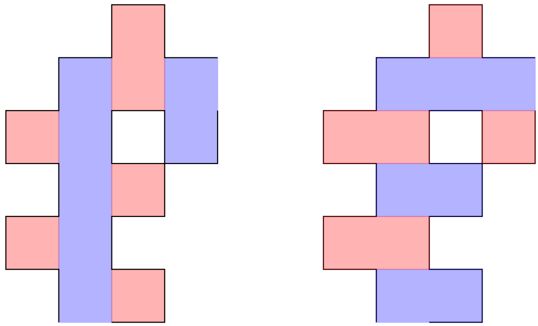

Figure 4.8: A section of a wire and its optimal tilings: vertical (left) and horizontal (right).

a wire uses either vertical or horizontal rectangles but not both. Any other tiling requires at least one more convex set.

The second tool is called a junction, and it serves to model a clause. Figure 4.9 depicts a junction corresponding to a clause of three literals. Figure 4.10 shows a zoom on the most important part of a junction. A junction has three arms, represented as dashed black line segments. A junction for a clause of two literals is obtained by blocking one of the arms of the junction. The blue line segments have length $1 + \varepsilon$, for a fixed ε arbitrarily small. Therefore, the red line segments are not aligned with the long black line segment to their right. A junction can be in four different orientations, which can be obtained successively by making rotations of 90°. Let us consider the orientation of the junction in Figure 4.9. One wire is connected from the left, one from above, and one from the right. A wire can only be connected to a junction at one of its bends (see Figure 4.7). We then remove the line segment corresponding to the arm of the junction, as illustrated in Figure 4.9.

If the tiling of the wire connected from the left is horizontal, then one of the rectangles can be prolonged into the junction. The same

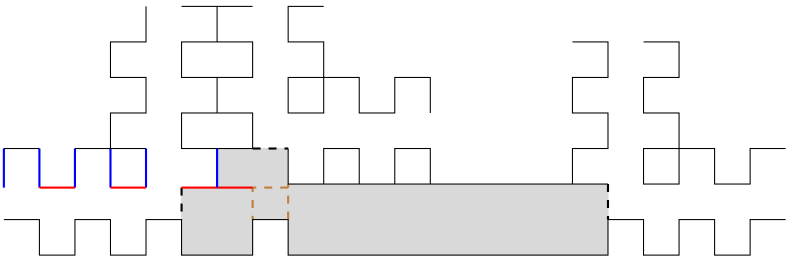

Figure 4.9: A junction for the MCT problem.

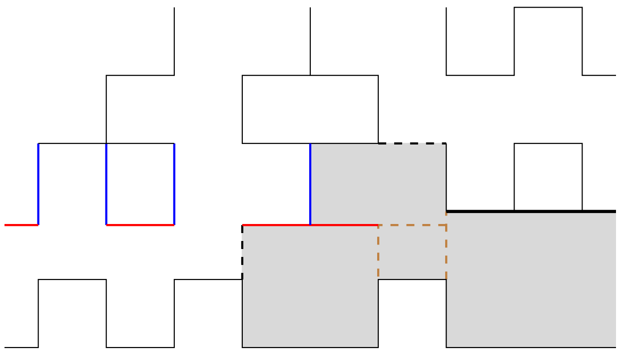

Figure 4.10: A zoom on a junction for the MCT problem. The blue segments have length $1 + \varepsilon$. One of the intersections between two dashed brown segments lies in the interior of the polygon. The red segments are not aligned with the long thick one to the right.

holds for the wire connected from the right. On the contrary, a rectangle can be prolonged from the wire connected from above only if the tiling is vertical. If a rectangle can be prolonged, we say that the wire *sends true*, otherwise it *sends false*. If a clause contains two negative literals $\bar{x}, \bar{y}$ and one positive z, the corresponding junction is as in Figure 4.9, or the 180° rotation of it. The wire corresponding to z is connected from above or below, and the wires corresponding to x and y are connected from the left and from the right, or vice versa. Therefore, the wire corresponding to z sends *true* if and only if z is set to *true*. On the contrary, the wire corresponding to x (respectively y) sends *true* if and only if x (respectively y) is set to *false*. If the clause has two positive literals, then the junction is vertical, and the junction behaves likewise.

Lingas proved that when minimising the number of convex polygons in a tiling, for each junction at least one adjacent wire sends *true*. Before stating Lingas' lemma exactly, we need to explain the first step of the construction of the point set.

4.4.1 Construction of the polygon with holes

Let us consider one instance (F, G) of MPLSAT. Lingas states that the planarity of G implies that the junctions and the wires can be embedded as explained above, and so that they do not overlap [58]. Thus we obtain a polygon with holes, that we denote by Π. He adds without proof that the dimensions of Π are polynomially related to $|V|$, where V denotes the vertex set of G. We show in the following paragraph how to embed the polygon with holes into a grid Λ, such that each edge consists of line segments of Λ. Actually, some parts of edges are not exactly line segments of Λ, but are shifted orthogonally by distance ε (recall the red segments in Figure 4.9). However, as ε is arbitrarily small, this does not impact our proof. Thus, for sake of simplicity, we will from now on do as if the edges entirely consisted of segments of Λ. Moreover, we show how to construct Λ in $\Theta(|V|^2)$

size.

Let us consider an instance (F, G) of MPLSAT. Recall that the vertex set of G is the union of X and C, where X denotes the set of variables and C the set of clauses. We define a new graph $G' = (V', E')$ as follow: For each vertex $x \in X$ of degree δ, we have δ vertices $x_1, \ldots, x_\delta$ in V'. Moreover for each $c \in C$ we have one vertex c in V'. Now for each edge $(x, c) \in E$, we have one edge between $c \in V'$ and one of the $x_i \in V'$. We do so such that each vertex x_i in V' corresponding to some $x \in V$ has degree 1. Then we add edges between the vertices $x_1, \ldots, x_\delta$ so that they induce a path. We can do this such that the graph we obtain, G', is still planar. Moreover, since a clause contains at most three vertices, the maximal degree of G' is at most 3. The number of vertices we have added is at most the number of edges in G, therefore $|V'| = \Theta(|V|)$. Following the result of Valiant [80], we can embed G' in a grid of size $\Theta(|V|^2)$, such that the edges consist of line segments of the grid. Let s be the line segment incident to c in the edge $\{x_i, c\}$, where x_i corresponds a variable $x \in X$ and c is a clause. Our constraint is that x appears positively in c if and only if s is vertical. Moreover we impose s to be of length at least 10, so that we have enough space later to replace c by a junction. We claim that we can find an embedding of G on another grid, still of size $\Theta(|V|^2)$, that satisfies our constraint. We first scale the embedding by 3. Then we can change the path of each edge adjacent to c as illustrated in Figure 4.11, so that the embedding satisfies our first constraint. The line segments in red (respectively blue) correspond to edges $\{x, c\}$ where x appears positively (respectively negatively) in c. By assumption, at most two variables appear positively, and at most two appear negatively. Therefore it is possible to adapt the paths of the edges so that the red line segments are vertical and the blue ones horizontal. Finally, we scale the grid by 10 to satisfy the second constraint.

Now we replace each clause by a junction. Let $x_1, \ldots, x_\delta$ be the

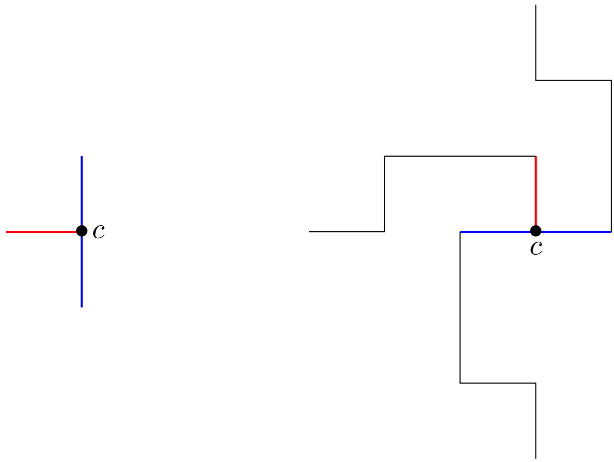

Figure 4.11: How to adapt the embedding to satisfy the constraint. On the left, before adaptation, on the right, after adaptation.

vertices corresponding to a variable $x \in X$. Observe that the set of line segments $\bigcup_{1 \leq i \leq \delta} \{e \in E' \mid x_i \in e\}$ is a tree. Let us take a constant sufficiently big, and scale the grid so that we can replace each tree corresponding to some variable $x \in X$ by a simple polygon that is as close as possible to the tree. Then, after a new scaling, we can replace each polygon by a wire, in a grid which is still of size $\Theta(|V|^2)$. As ε introduced in the definition of the junctions is arbitrarily small, we can consider that all line segments of junctions and wires are segments of a grid. For sake of simplicity, we will from now on omit to mention that some line segments are not exactly on the grid. Again, it is possible to find an embedding in a grid of size $\Theta(|V|^2)$, such that wires do not intersect, and a wire is connected to a junction if and only the variable corresponding to the wire is contained in the clause corresponding to the junction.

We can now state Lingas' lemma:

Lemma 4.26 (Lingas [58])**.** *In a minimum tiling with convex sets of Π, a junction contains wholly at least three convex sets. The junction contains wholly exactly three if and only if at least one of the wires connected to the junction sends* true*.*

4.4.2 Discretisation of the line segments

To construct the point set of the reduction from MPLSAT to MCT, we first construct a collection of line segments. We then discretise this collection by replacing each line segment by a set of collinear points.

Let us consider our polygon with holes Π that lies in the grid Λ. The grid consists of points with integer coordinates, and line segments between points that are at distance 1. We consider the collection of line segments consisting of Π union each line segment of Λ whose relative interior is not contained in the interior of Π. Notice that therefore we have line segments outside Π, but also inside its holes. Moreover, the collection of line segments that we obtain, denoted by Φ, is a subgraph of the grid graph Λ.

Now we define K as twice the number of unit squares in Λ plus 1. Finally, we replace each line segment in Φ by K points evenly spaced. We denote this point set by P.

4.4.3 Proof of correctness

We have constructed P in order to have the following property:

Lemma 4.27. *In a minimum convex tiling Σ of P, for each convex set $S \in \Sigma$, the interior of S does not intersect Φ.*

Before proving Lemma 4.27, we first explain how we use it. Let K' denote the number of unit squares in Φ, plus the minimum number of rectangles in a partition of the wires, plus three times the number of clauses. Using Lemmas 4.25 and 4.26 shown by Lingas coupled with Lemma 4.27, we immediately obtain the following theorem:

Theorem 4.28. *The formula F is satisfiable if and only if there exists a convex tiling of P with K' polygons.*

Since P and K' can be computed in polynomial time, Theorem 4.28 implies Theorem 4.24 for the MCT problem. We use a packing argument, and claim that in a convex tiling Σ of P, if a convex set $S \in \Sigma$ has large area, then most of its area is contained in a unique cell of Φ. Then we show that in a minimum convex tiling, all convex sets have large area, and that each of them fills the cell that contains it. For a set S, let $\mathrm{A}(S)$ be the area of S.

Lemma 4.29. *Let L and L' be two squares in Λ, and S be a convex polygon whose interior does not contain any point in P. If $\mathrm{A}(S \cap L) > 1/K$, and the boundary of S crosses a line segment of Φ between L and L', then $\mathrm{A}(S \cap L') \leq 1/K$.*

Proof. The proof is illustrated in Figure 4.12. By assumption, S intersects a line segment whose endpoints p and q are at distance $1/K$. Let us consider the two line segments s and s' of the boundary of S that intersect the line ℓ which contains p and q. Assume for contradiction that the lines containing respectively s and s' do not intersect, or intersect on the side of ℓ where L lies. This implies that $S \cap L$ is contained in a parallelogram that has area $1/K$, as illustrated in Figure 4.13. Indeed such a parallelogram has base $1/K$ and height 1, therefore $\mathrm{A}(S \cap L) \leq 1/K$. This shows that the lines containing respectively s and s' intersect on the side of ℓ where L' lies. Using the same arguments as above, this implies $\mathrm{A}(S \cap L') \leq 1/K$. □

Lemma 4.30. *Let R be a rectilinear polygon on Λ whose interior does not contain any point in P. Let Ψ be a set of squares of Λ contained in R, such that no two of them are on the same row or*

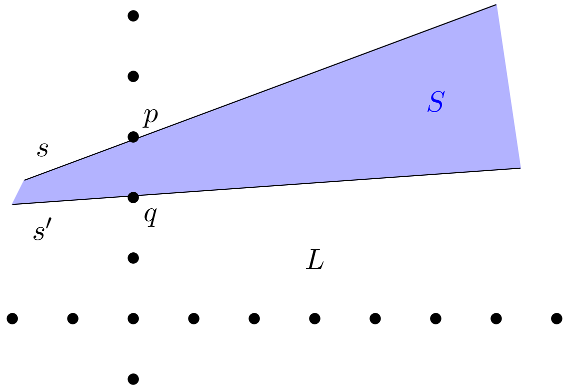

Figure 4.12: If $A(S \cap L) > 1/K$, the two lines containing s and s' intersect on the left side.

same column of Λ. Let Σ be a minimum convex tiling of P. Let S be an element of Σ, such that for each $L_i \in \Psi$, we have $\mathrm{A}(S \cap L_i) > 1/K$. Then there exist $|\Psi| - 1$ squares $\{L_i\}$ in Ψ and $|\Psi| - 1$ convex sets $\{S_i\}$ in Σ, such that for any i: $\mathrm{A}(S_i \cap L_i) > 1/K$.

Proof. We assume $|\Psi| \geq 2$, otherwise there is nothing to prove. Using Lemma 4.29, we know that if such a S exists, then there are no line segment of Φ between any two squares in Ψ. We can observe thanks to how wires and junctions are defined that $|\Psi|$ is at most three. Moreover for the same reason, there are $|\Psi| - 1$ squares in Ψ such that the area of their intersection with S is at most $1/2$, as illustrated in Figure 4.14. Observe that by taking each unit square in Λ as a convex face, we obtain a convex partition with $\frac{K-1}{2}$ convex faces. In particular, there are at most K convex faces in Σ, which implies that for each unit square $L \in \Psi$, there exists a convex face $\tilde{S} \in \Sigma$ with $\mathrm{A}(\tilde{S} \cap L) > 1/K$. Moreover by iterating this process we can choose these convex sets so that they are distinct. □

Proof of Lemma 4.27. We define a function $f_0 : \Lambda \to \Sigma$ that maps

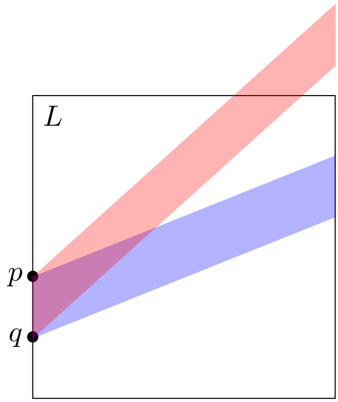

Figure 4.13: The area of the parallelograms is $1/K$.

each square L in Λ to a convex set $S \in \Sigma$ such that $\mathrm{A}(S \cap L) > 1/K$. Such a set exists since otherwise we would need more than K sets only to fill the square L. Now we define a map f in several steps. We start from f_0. At step $i+1$, find L and L' on the same line or the same column, and L'' between them, such that $f_i(L) = f_i(L') \neq f_i(L'')$. If this is not possible then stop. Otherwise define f_{i+1} as identically equal to f_i, except for all square between L and L', that are mapped to $f_i(L)$. Notice that f_{i+1} keeps the property that for each square $\tilde{L}$, we have $\mathrm{A}(f_i(\tilde{L}) \cap \tilde{L}) > 1/K$. Moreover, this procedure must stop eventually because the interiors of the convex sets are non-overlapping. We denote by f the map obtained after the last iteration.

We denote by f^{-1} the function that maps a set $S \in \Sigma$ to $\{L \in \Lambda \mid f(L) = S\}$. By Lemmas 4.29 and 4.30, if for some S there are m squares in $f^{-1}(S) \geq 2$ such that no pair of squares are on a line nor on a column, then there exists at least $m-1$ convex sets S_i such that for each S_i, $f^{-1}(S_i) = \emptyset$. Moreover, such a S_i cannot appear when considering another element of Σ: If Lemma 4.30 when applied to S gives the existence of the $\{S_i\}$, and when applied to $S' \neq S$ gives

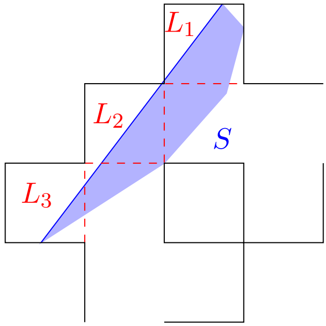

Figure 4.14: If $A(S \cap L_1) \geq 1/2$, then $A(S \cap L_2) \leq 1/2$ and $A(S \cap L_3) \leq 1/2$.

the existence of the $\{S'_j\}$, then all these convex sets are distinct. We denote by M the number of such sets, taken over all sets $S \in \Sigma$ with $f^{-1}(S) \geq 2$.

We are going to partition P into at most $|\Sigma| - M$ rectilinear polygons (not necessarily convex). We denote by T the partition. For each set $S \in \Sigma$ such that $|f^{-1}(S)| = 1$, we add in T the square in $f^{-1}(S)$. For each $S \in \Sigma$ such that $|f^{-1}(S)| \geq 1$, we add in T the rectilinear polygon consisting of the union of the squares in $f^{-1}(S)$. We know thanks to Lemma 4.29 that the boundary of this rectilinear polygon does not cross any line segment in Φ. Following what was argued above, there are at most $|\Sigma| - M$ sets in T. We now construct a new convex tiling Σ', that we claim to be minimum. We replace each set in T that is not convex by rectangles. From the construction of wires and junctions, there exists a partition of the sets in T into rectangles that will add M new sets. Looking at the construction of T and then Σ', we can observe that the convex tiling Σ was minimum if and only if the interior of all sets S does not intersect Φ. $\square$

4.4.4 How to adapt the proof to MCP

Let us explain how to adapt the proof to the case of the MCP problem. Notice that the only lemma for which it makes a difference is Lemma 4.26. Indeed, if in Figure 4.9 the wire connected from below sends *true*, then one can prolong a rectangle of this wire so that it contains the small rectangle in the junction. Therefore only three rectangles are wholly contained in the junction. However, the vertex on the top left of the prolonged rectangle is not a point of P. This can be solved by adding a line segment to the construction, coloured in red in Figure 4.15. With this construction, the rectangle is prolonged into a trapezoid (delimited above by the dashed red line segment). The convex set above remains a rectangle, and the one to the left becomes a convex quadrilateral with a right angle. Now one can adapt the proof by taking into account this new triangular hole that has been inserted in each junction. The proof can then be done similarly to the one of the MCT problem.

4.4.5 NP-hardness of CPNCS

Let us consider the CPNCS problem with additional constraints on the input point set P. We consider the CPNCS problem on a 3-directed point set P with set of directions L, such that there are no five collinear points, and a point $p \in P$ is contained in at most one line in L. We show that CPNCS is NP-hard, even for such constrained point sets.

Kratochvíl and Nešetřil have shown that *Maximum Independent Set in Intersection Graphs of Segments* is NP-hard, even when the segments lie in only three possible directions [53]. The problems corresponds to finding a maximum set of pairwise non-intersecting segments. There is an additional constraint that no two parallel segments intersect. We do a reduction from this problem to CPNCS with the additional constraints. The reductions is illustrated in Figure 4.16. Let us consider a set of such line segments. We first shift

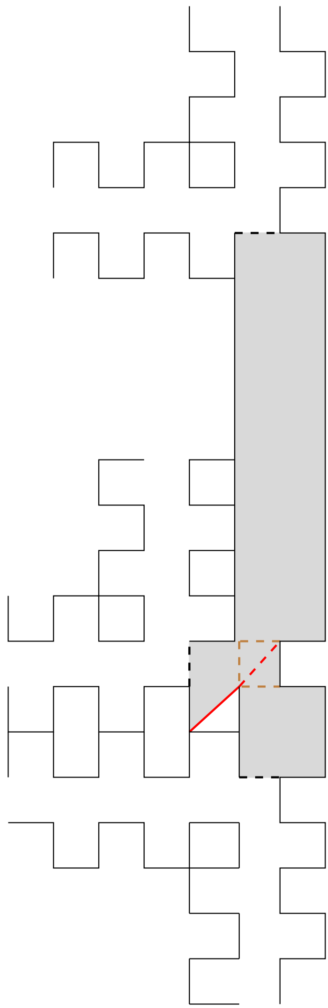

Figure 4.15: A junction for the MCP problem.

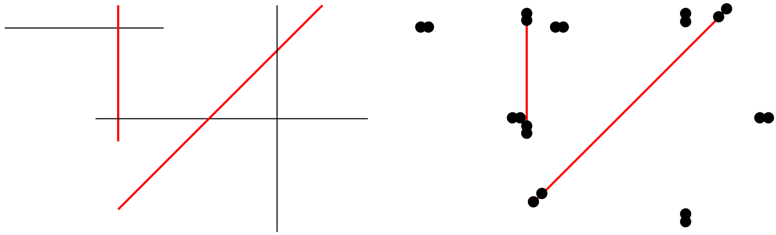

Figure 4.16: The reduction from Maximum independent set in segments to CPNCS. A segment in an independent set corresponds to a segment in a covering which covers four points instead of only two.

the segments so that no two segments are aligned. Since no two parallel segments intersect, this does not change the intersection graph. We also extend some segments such that if two segments intersect, then the intersection is in the relative interiors of the segments. We replace each segment by four collinear points, that lie on the segment. Let us consider a segment s with endpoints u and v. By construction, there is a connected subset of s that contains u but no intersection point between s and another segment. We add two points in this connected subset. The same holds with v, and we add two other points in the corresponding connected subset of s that contains v. We do it such that three points are collinear only if they lie on the same segment of the input set. It is clear that the point set satisfies our constraints.

We denote by n the number of segments. We claim that there is an independent set of s segments if and only if there is a covering of the points with $2n - s$ non-crossing segments. Assume there is an independent set of s segments, and let us show the existence of a covering with $2n - s$ elements. For each segment in the independent set, we take it in the covering. For the other segments, we take two segments to cover the four points: one for each pair of close points. Therefore we have covered the points with $s + 2(n - s)$ segments,

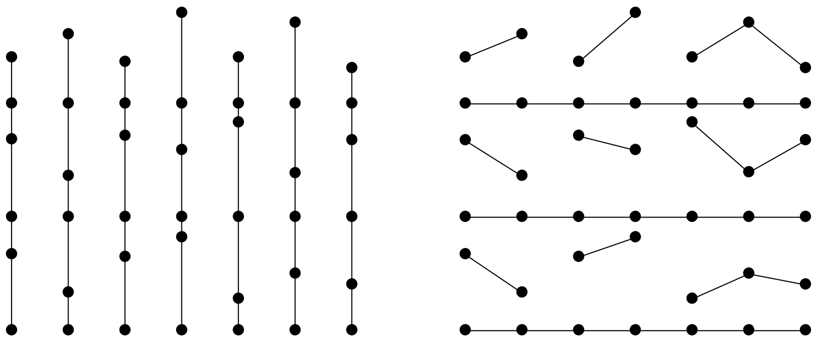

Figure 4.17: On the left a minimum covering of the point set. On the right a covering given by the greedy algorithm. As the horizontal segments contain seven points, they are the firsts taken by the greedy algorithm. However, afterwards many segments are needed to cover the remaining points.

which are non-crossing by assumption.

Let us now assume that there is a covering of the point set with $2n - s$ non-crossing segments. Let us denote by x the number of segments that cover more than two points. Such a segment covers at most four points, and there are $4n$ points to cover. Hence, we have $4x + 2(2n - s - x) = 4n$, which implies $x = s$. By construction, each segment that covers more than two points is a segment of the input set, thus there exists an independent set of s segments.

4.4.6 Lower bound for the greedy algorithm

The construction is illustrated in Figure 4.17. Consider the set of $2k(2k+1)$ points $P := \{(x, y) \mid x, y \in \mathbb{N}, 1 \leq x \leq 2k+1, 1 \leq y \leq 2k\}$. Now we perturb slightly the points such that any line that covers at least three points is either vertical or horizontal. Actually, after perturbation, the lines might not be exactly vertical or horizontal, but we keep referring to them as such for simplicity. Moreover, we perturb again the points such that only k horizontal lines cover at least three points. We do so such that, for any pair of consecutive horizontal lines that covers at least three points, there are $k+1$ points in between (which correspond to the points of the former horizontal line). In the end, a vertical line covers at least three points if and only if it covers exactly $2k$ points. Similarly, a horizontal line covers at least three points if and only if it covers exactly $2k+1$ points. We denote this new point set by P', and denote by n the cardinality of P'. It is clear that one can cover all points using $2k+1$ vertical segments (which are thus non-crossing). However, the greedy algorithm will first pick the k horizontal segments with even y-coordinate. Then it remains to cover the $k(2k+1) = \frac{n}{2}$ remaining points, but there is no line covering three of these points. Therefore the greedy algorithm picks at least $\frac{k(2k+1)}{2}$ more segments. Hence, the approximation ratio is at least $\frac{k(2k+1)/2}{2k+1} = k/2 = \Omega(\sqrt{n})$.

Recall that in Theorem 4.5, we introduce another algorithm whose approximation ratio is $\mathcal{O}(\sqrt{n}\log n)$ for solving CPNCS. The algorithm first computes a minimum covering with lines of the point set, and then divides these lines into non-crossing segments. By the same reasoning, this algorithm will also have an approximation ratio in $\Omega(\sqrt{n})$ on the point set P'.

CHAPTER 5

Triangulations and the bistellar flip graph

5.1 Introduction

In Chapter 4 we studied convex partitions of points sets in degenerate position. A convex partition where all faces are triangles is a "triangulation". In this chapter, we study triangulations on degenerate point sets. Triangulations provide structure on point sets, flip graphs provide structure on the set of all triangulations of a given point set, see e.g. [22]. A triangulation can be flipped to another if they are sufficiently similar. As we detail later more formally, this notion of flip is actually connected to the concept of convex partitions.

Since the seventies, it has been known that in the plane these flip graphs are connected, [56], while it is known since the turn of the century that there are point sets in dimension 5 or higher for which this graph (more precisely, the so-called bistellar flip graph) is not connected, [72, 73]; the situation in dimension 3 and 4 remains open. De Loera, Rambau and Santos, [22], raised the question of the vertex-connectivity of the flip graph in the plane, which was recently answered for point sets in general position (i.e. no three points on a line), [82]. The goal of this chapter is to complete the picture in the plane by showing that the bistellar flip graph of point sets in the plane is always $(n-3)$-vertex-connected. That is, we allow sets in which three or more points lie on a common line. Following [22], we include even the possibility of repeated points, a situation that is modelled by so-called *point configurations* (to be defined below).

A motivation for considering such point configurations with repeated points comes, e.g. from the fact that in some applications the points we are considering are projected from a higher dimension where some points may have the same image. Other motivations come from Gale duality, see [22, Sec. 2.1.2].

For the rest of this section, we provide the formal setting for the chapter (as given in [22]) and the statement of the results.

5.1.1 Point configurations and triangulations

Formally, we define a *point configuration* A as a set of labels J with a mapping $J \to \mathbb{R}^2$; we write A as $(p_j)_{j \in J}$. A *subconfiguration* $B = (p_j)_{j \in I}$ of A induced by a subset I of the labels J is called affinely independent if the points labelled by I are affinely independent – in particular, they are all distinct. We denote by *conv*(B) the convex hull of the points $\{p_j\}_{j \in I}$ with labels in I, and we denote by *relint*(B) the relative interior of *conv*(B).

While triangulations of point sets in general position can be defined as maximal straight line embedded plane graphs, see [82] for example, the situation is more subtle for point configurations. The

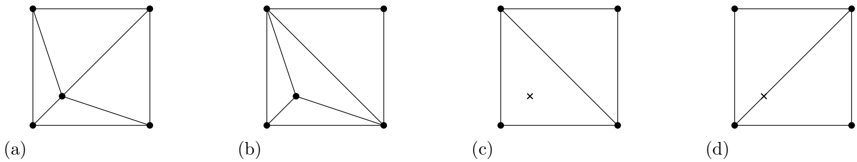

Figure 5.1: The four triangulations of a point configuration with no repeated point. A cross indicates a skipped point.

following is the established definition (as a simplicial complex) for this set-up, see [22].

Definition 5.1 (Triangulation)**.** A *triangulation* of a point configuration A is a collection $\mathcal{T}$ of affinely independent subconfigurations with the following properties:

1. If B is in $\mathcal{T}$ and F is a subset of B, then F is in $\mathcal{T}$. (Closure Property)
2. If B and B' are two different elements in $\mathcal{T}$, then $relint(B) \cap relint(B') = \emptyset$. (Intersection Property)
3. The union $\bigcup_{B \in \mathcal{T}} conv(B)$ equals $conv(A)$. (Union Property)

A set of one, two, or three elements in $\mathcal{T}$ is called a *vertex*, an *edge*, or a *triangle*, respectively, of the triangulation $\mathcal{T}$. Note that no set in $\mathcal{T}$ contains more than three elements because of affine independence. Observe that some labels may not appear in any set of $\mathcal{T}$ – such elements are called *skipped* in $\mathcal{T}$. Figure 5.1 shows the four triangulations of a point configuration without repeated points. In triangulations (c) and (d) one point is skipped. Observe that this skipped point might lie on a triangle, see (c), or on an edge, see (d). It might also lie on an edge on the boundary of the convex hull of the point configuration. A skipped point may even lie on a vertex in the case of repeated points.

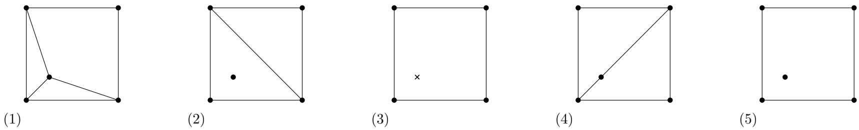

Figure 5.2: The five subdivisions which are not triangulations of a point configuration with no repeated point. A cross indicates a skipped point.

5.1.2 Flips: Via faces, subdivisions, and coarsenings

In order to proceed to the definition of bistellar flip graphs of point configurations, we need to introduce flips, the minimal changes applied to a triangulation. Point and edge flips allow for a simple direct definition in general position, see [82], for example. Here we need to first introduce faces and (polyhedral) subdivisions.

We start off with the definition of a face of a point configuration, as given in [22]. These are subconfigurations which can be obtained as intersections of the point configuration with a face (in the "conventional sense" as a convex set) of the convex hull of the point configuration. Or, equivalently, these are the subconfigurations of all labels whose points maximize some given linear function.

Definition 5.2 (Face)**.** Let $A = (p_j)_{j \in J}$ be a point configuration. Let C be a subset of J. For a linear function $\psi : \mathbb{R}^2 \to \mathbb{R}$, the face of C in direction ψ is the set $\mathit{face}_A(C, \psi) := \{j \in C \mid \psi(\mathbf{p}_j) = \max_{b \in C}(\psi(\mathbf{p}_b))\}$.

We say that F is a *face of* C if it is a face for some linear function ψ. Observe that C is always a face of C as obtained with the zero function. Moreover, we consider the empty set as a face of C. We are now ready to define subdivisions as in [22].

Definition 5.3 (Subdivisions)**.** A *subdivision* of a point configuration $A = (p_j)_{j \in J}$ is a collection $\mathcal{S}$ of subsets of J that satisfies the following properties:

1. If B is in $\mathcal{S}$ and F is a face of B, then F is in $\mathcal{S}$. (Closure Property)
2. If B and B' are two different elements in $\mathcal{S}$, then $relint(B) \cap relint(B') = \emptyset$ (Intersection Property)
3. The union $\bigcup_{B \in \mathcal{S}} conv(B)$ equals $conv(A)$. (Union Property)

The elements of a subdivision are called *cells*. As before for triangulations we say that a label that does not appear in any set of $\mathcal{S}$ is *skipped*, and that a cell with one element is called a *vertex*. A cell whose convex hull has dimension k is called a *k-cell*. Roughly speaking, subdivisions look like convex partitions as defined in Chapter 4. Recall that convex partitions are geometric plane graphs. According to our definition, subdivisions are not. However, for each subdivision we can associate a geometric plane graph. For each 1-cell we draw a segment between the two extreme points of that 1-cell. Figure 5.2 shows the five subdivisions which are not triangulations of the point configuration from Figure 5.1. Observe that a cell may contain more than two elements and still be a 1-cell, as in subdivision (4). Likewise, as we allow point repetitions, a cell may contain arbitrarily many labels and still be a 0-cell.

The following remark lists some differences between convex partitions and geometric plane graphs associated with subdivisions.

Remark 5.4. In subdivisions some points may be skipped, but are still taken into consideration. Moreover, there might be isolated points. Some of them are in some 2-cells, and some lie on 1-cells. They may look like vertices of degree 2, but it is important to keep in mind that they are actually just isolated points (of degree 0) that lie on some 1-cell. Finally, if a point is not extreme in one of the convex faces, then it has degree 0. This is shown in Figure 5.3. If

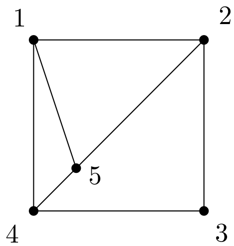

Figure 5.3: A convex partition which is not a subdivision (see Remark 5.4). If it were a subdivision $\mathcal{S}$, then $\mathcal{S}$ would be equal to $\{1,4,5\}$, $\{1,2,5\}$, $\{2,3,4,5\}$ and all faces of these sets. However, the intersection of the interior of the faces $\mathit{relint}(\{2,5\}) \cap \mathit{relint}(\{2,4,5\})$ is not empty.

$\mathcal{T}$ were a subdivision, the 2-cells in $\mathcal{T}$ would be $\{1,5,4\}$, $\{1,2,5\}$ and $\{2,3,4,5\}$. But this is not valid as therefore $\{2,5\}$ and $\{2,4,5\}$ would be 1-cells of $\mathcal{T}$ (closure property), but then $\mathit{relint}(\{2,5\}) \cap \mathit{relint}(\{2,4,5\})$ is not empty, in contradiction with the intersection property.

Definition 5.5 (Refinement and coarsening)**.** Given two subdivisions $\mathcal{S}$ and $\mathcal{S}'$, $\mathcal{S}$ is a *refinement* of $\mathcal{S}'$ if for each cell B of $\mathcal{S}$, B is a subset of a cell B' of $\mathcal{S}'$. We also say that $\mathcal{S}'$ is a *coarsening* of $\mathcal{S}$.

The coarsening relation gives a partial order. The minimal elements are the triangulations, which cannot be refined. There is a unique maximal element, called the *trivial subdivision* $\mathsf{S}_{\mathsf{triv}}$, where all faces of A are in $\mathsf{S}_{\mathsf{triv}}$ (including J itself). Figure 5.4 shows the Hasse diagram of the coarsening relation for the point configuration of Figures 5.1 and 5.2.

Definition 5.6 (Height)**.** *The height of a subdivision $\mathcal{S}$ in the partial order given by the coarsening relation is equal to the size of a longest chain in the Hasse diagram ending in $\mathcal{S}$ minus one.*

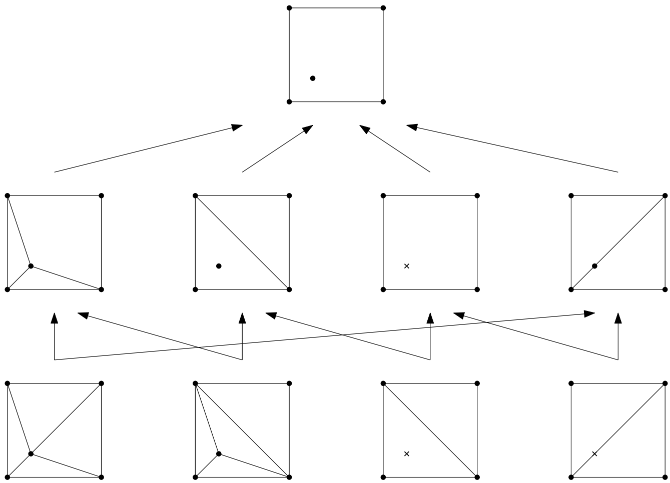

Figure 5.4: Hasse diagram of a point configuration where no point is repeated. The partial order displayed corresponds to coarsenings.

Note that the height of a triangulation is 0. We also give a special name to subdivisions of height 1.

Definition 5.7 (Almost-triangulations)**.** A subdivision at height one in the poset of subdivisions is called an *almost-triangulation*. That is, it is not a triangulation but all its proper refinements are triangulations.

Lemma 5.8 ([22])**.** *Any almost-triangulation has exactly two proper refinements, which are both triangulations.*

Definition 5.9 (Flips [22])**.** We say that two triangulations of the same point configuration are connected by a flip supported on the almost-triangulation $\mathcal{S}$ if they are the two triangulations refining $\mathcal{S}$.

As noted in [22], every flip happens on a *circuit*, where a circuit is a minimal affinely dependent subset of points, and splits in a unique way into a pair (Z_+, Z_-) with the property that $conv(Z_+) \cap conv(Z_-) \neq \emptyset$. We recall the definition of a circuit to explain the names of the types of flips. We say that (Z_+, Z_-) is a circuit of *type* $(|Z_+|, |Z_-|)$, with the convention $|Z_+| \geq |Z_-|$.

Let us consider a point configuration in the plane. There are four types of flips [22]. The first is the flip of type $(2, 2)$. It occurs when there is a 2-cell with four elements in the almost-triangulation $\mathcal{S}$, such that the four corresponding points are in convex position. This 2-cell is the only non-simplicial cell of $\mathcal{S}$. The two triangulations are obtained by adding one or the other diagonal. This can be seen in Figure 5.1: one can go from (a) to (b) and vice versa with an edge flip. The corresponding almost-triangulation is subdivision (1) in Figure 5.2. Likewise, one can go from (c) to (d) and vice versa in Figure 5.1, as those are the two proper refinements of subdivision (3) in Figure 5.2.

The second flip is the $(3, 1)$-flip. It occurs when there is a 2-cell with four elements in the almost-triangulation $\mathcal{S}$, such that one point p is in the convex hull of the three others. The two triangulations are obtained by either skipping p, or by connecting it to the three other points. This operation can be seen in Figure 5.1 when going from (b) to (c) and reciprocally. Those two triangulations are the proper refinements of subdivision (2) in Figure 5.2.

The third flip is the $(2, 1)$-flip. It occurs when there are three collinear points in one or two 2-cells (depending on whether they lie on the boundary of the convex hull or not). Let us denote by p the point in between the two others. In this flip, we remove the edge on which p lies, and we connect p to all vertices of the triangles that were adjacent to the edge. This point has now degree 4 if it is an inner point, and degree 3 if it lies on the boundary of the convex hull. The converse operation consists in finding a point adjacent to four if it is an inner point, or to three if it lies on the boundary of the convex hull, such that two of the incident edges are aligned. The point becomes skipped, the edges are removed, and a new edge covering the two that were aligned is added. This flip can be seen in Figure 5.1 when going from (a) to (d). Those two triangulations are the proper refinements of subdivision (4) in Figure 5.2.

The fourth flip is the $(1, 1)$-flip. When a point is repeated, with one label j not skipped and the other j' being skipped, it consists in removing the edges incident to j and making them incident to j'. Also, j becomes skipped and j' becomes not skipped. This has no visual effect on the triangulation.

Definition 5.10 (Bistellar flip graph)**.** The vertex set of the *bistellar flip graph* of a point configuration A is the set of all triangulations on A. Two triangulations are adjacent if there is an almost-triangulation that coarsens both.

The bistellar flip graph of the point configuration depicted in Figure 5.1 consists of four vertices. Observe that this graph is a cycle of length 4.

In [22], De Loera, Rambau and Santos asked whether for all point configurations of size $n \geq 3$ in the plane, the bistellar flip graph is $(n-3)$-connected. Wagner and Welzl showed that this is true when the configuration is in general position [82], meaning that no three points are on a line, thereby also implying that no point is repeated. Observe that therefore they did not consider $(2, 1)$-flips and $(1, 1)$-flips. Allowing for points to be repeated, or for points to be on a line, does not only create new kinds of vertices in the bistellar flip graph. It also creates new types of edges, corresponding to those new types of flips. We extend their result to all point configurations in the plane.

Theorem 5.11. *The bistellar flip graph of a point configuration with convex hull of dimension d, where $d \leq 2$, with n labels is $(n-d-1)$-*

connected, and this is tight for all point configurations.

We immediately prove the result for a point configuration A whose convex hull has dimension 0. Observe that in a triangulation of A all points but one are skipped. In an almost-triangulation, all points but two are skipped. The bistellar flip graph is isomorphic to a complete graph of size n, and is therefore $(n-1)$-connected but not n-connected.

Let us now consider a point configuration A whose convex hull is 1-dimensional. In a triangulation of A, every point that is skipped can be flipped via a $(2,1)$-flip or a $(1,1)$-flip. Every point that is not skipped can be flipped via a $(2,1)$-flip, except for the two extreme points. Therefore, the bistellar flip graph is $(n-2)$-regular, and cannot be $(n-1)$-connected. If no point is repeated, then the bistellar flip graph is the hypercube graph of dimension $n-2$, which is $(n-2)$-connected.

Finally, let us consider a point configuration A whose convex hull is 2-dimensional. As shown in [23], there is a triangulation that has degree $n-3$ in the bistellar flip graph of A, and all triangulations have degree at least $n-3$. We show the following:

Theorem 5.12. *The bistellar flip graph of a point configuration in the plane where no point is repeated with $n \geq 3$ points is $(n-3)$-connected.*

Assuming that no point is repeated simplifies quite a lot the definitions. This is why we prove the theorem for point configurations where no point is repeated, which can be seen as a point set, and we extend it to all point configurations in a separate step. We have argued above that the bistellar flip graph of a point configuration with 1-dimensional convex hull where no point is repeated is $(n-2)$-connected. Therefore, assuming that Theorem 5.12 holds, to prove Theorem 5.11 it remains to show that if it holds for point configurations where x points are repeated, then it also holds for point configurations where $x+1$ points are repeated. This is what we do in the next section.

5.2 From point sets to point configurations

In this part, we assume that Theorem 5.12 holds, and we show by induction on the number of repeated points that Theorem 5.11 holds too. If we consider a point configuration with no repeated points, then we are immediately done. Let us now assume that Theorem 5.11 holds for all configurations where at most x points are repeated. We show that the theorem holds for all configurations where $x+1$ points are repeated. Let A' be a configuration where $x+1$ points are repeated, including point p. We denote by ℓ the number of labels which are mapped to p in A'. Let A denote the point configuration identical to A', except for point p which is not repeated. We want to describe the relation between the bistellar flip graphs of A and A'. To do so, we first define a concept.

Let $G = (V, E)$ be a graph. Given an induced subgraph H of G, and any positive integer ℓ, we define below the (H, ℓ)-*clone* of G, denoted by G'. First, let us give an intuitive definition. The graph G' is the union of ℓ graphs isomorphic to H, union a graph isomorphic to $G \setminus H$, with few more edges. The corresponding vertices of two different copies of H are connected by an edge. Moreover, when considering the union of two graphs, one isomorphic to H and the other to $G \setminus H$, this union is isomorphic to G.

We define now the (H, ℓ)-clone formally. It is depicted on Figure 5.5. Let n and m denote the number of vertices in G and H, respectively. We denote by $\{v_j\}_{1 \leq j \leq m}$ the vertices in H, and by $\{v_j\}_{m+1 \leq j \leq n}$ the vertices in $G \setminus H$. We define a set of ℓm vertices, denoted by $v_{i,j}$, for $1 \leq i \leq \ell$ and $1 \leq j \leq m$. We denote by V_i the set $\{v_{i,j} \mid 1 \leq j \leq m\}$, for all $1 \leq i \leq \ell$. We also define $V_0 := \{v_{0,j} \mid m+1 \leq j \leq n\}$ a

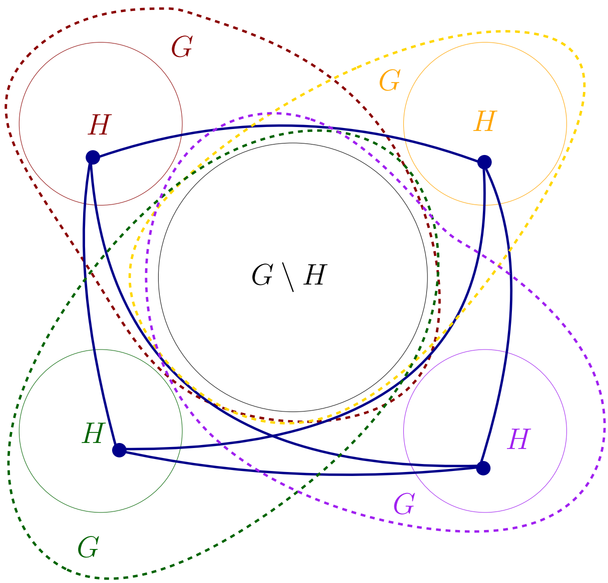

Figure 5.5: A $(H, 4)$-clone of a graph G. The union of $G \setminus H$ and a copy of H is isomorphic to G. Moreover, the corresponding vertices of two different copies of H are connected by an edge.

set of $n-m$ vertices. We define V' as $\cup_{0\leq i\leq \ell} V_i$. Finally, we define E' by the four following properties. Property 1. states that the subgraph induced by V_0 is isomorphic to $G \setminus H$, Property 2. that for all $1 \leq i \leq \ell$, the subgraph induced by $V_0 \cup V_i$ is isomorphic to G, Property 3. that for all $1 \leq i \leq \ell$, the subgraph induced by V_i is isomorphic to H and Property 4. states that the corresponding vertices of two different copies of H are connected by an edge.

1. for all $m+1 \leq j_1, j_2 \leq n$, we have $\{v_{0,j_1}, v_{0,j_2}\} \in E' \iff \{v_{j_1}, v_{j_2}\} \in E$,
2. for all $1 \leq i \leq \ell$, $1 \leq j_1 \leq m$, $m+1 \leq j_2 \leq n$, we have $\{v_{i,j_1}, v_{0,j_2}\} \in E' \iff \{v_{j_1}, v_{j_2}\} \in E$,
3. for all $1 \leq i \leq \ell$, $1 \leq j_1, j_2 \leq m$, we have $\{v_{i,j_1}, v_{i,j_2}\} \in E' \iff \{v_{j_1}, v_{j_2}\} \in E$,
4. for all $1, \leq i_1, i_2 \leq \ell$, $i_1 \neq i_2$, $1 \leq j_1, j_2 \leq m$ we have $\{v_{i_1,j_1}, v_{i_2,j_2}\} \in E' \iff j_1 = j_2$.

Finally, we define G' as the graph (V', E'). Recall that we have defined A to be a point configuration where a point p is not repeated, and a point configuration A' where p is repeated ℓ times. Let us denote by $G = (V, E)$ the bistellar flip graph of A, and by $G' = (V', E')$ the bistellar flip graph of A'. Let us denote by V_H the set of vertices in V corresponding to triangulations of A where p is not skipped. Observe that for all triangulations in A where p is not skipped, there are ℓ triangulations in A' where one of the repetitions of p is not skipped. Those are the only new triangulations that we can have by allowing the repetition of p. Therefore, V' is the union of $V \setminus V_H$ and ℓ copies of V_H. Moreover, the subgraph of G' induced by $V \setminus V_H$ union a copy of V_H is isomorphic to G. Finally, two triangulations from different copies of V_H are connected by an edge in G', as one can go from one to the other via a $(1,1)$-flip. Let us denote by H the graph induced by V_H in G. Therefore, G' is isomorphic to the (H, ℓ)-clone of G. Let us denote by n the number of labels in A. Thus, the number of labels in A' is equal to $n+\ell-1$. We are assuming that G is $(n-3)$ connected. We have to show that G' is $(n+\ell-4)$-connected. To do this we use the following lemma:

Lemma 5.13. *Let G be a graph, and let H be an induced subgraph of G such that any vertex v in $G \setminus H$ is connected to exactly one vertex in H. If G is k-connected, H is connected, and $G \setminus H$ is $(k-1)$-connected, then for any positive integer ℓ, the (H, ℓ)-clone of G is $(k+\ell-1)$-connected.*

Before showing Lemma 5.13, let us explain how it implies that Theorem 5.11 holds for all configurations where $x+1$ points are repeated. We only have to argue why any vertex v in $G \setminus H$ is connected to exactly one vertex in H, why G is $(n-3)$-connected, why H is connected, and why $G \setminus H$ is $(n-4)$-connected.

First, let us consider the bistellar flip graphs of A and the one of $A \setminus \{p\}$. Observe that the bistellar flip graph of $A \setminus \{p\}$ is isomorphic to the subgraph of G induced by $G \setminus H$. For both of them, x points are repeated, so we can apply induction. There are n labels in A, and therefore $(n-1)$ labels in $A \setminus \{p\}$ because p is not repeated in A. This implies that G is $(n-3)$-connected and that $G \setminus H$ is $(n-4)$-connected. We claim that H is connected. Take two triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ in H. If the two triangulations use different repetitions of a point q, we can flip one of the triangulations so that they use the same label of this repeated point. Let us now assume that if $\mathcal{T}_1$ and $\mathcal{T}_2$ use the same point, then it is with the same label. Let $\mathcal{P}_1$, respectively $\mathcal{P}_2$, denote the set of points that are not skipped in $\mathcal{T}_1$, respectively $\mathcal{T}_2$. Observe that it is always possible to flip a point which is currently skipped, either by a $(3,1)$-flip or by a $(2,1)$-flip. So we can flip in $\mathcal{T}_1$ the skipped points in $\mathcal{P}_2 \setminus \mathcal{P}_1$ to obtain a triangulation $\mathcal{T}_1'$ with set of points being equal to $\mathcal{P}_1 \cup \mathcal{P}_2$. We do the same for $\mathcal{T}_2$ and obtain $\mathcal{T}_2'$. Now by Lawson's Theorem [56] we can go from $\mathcal{T}_1'$ to $\mathcal{T}_2'$ by doing $(2,2)$-flips. Observe that during this process p is never skipped. Thus H is connected. Finally, we claim that any vertex v in $G \setminus H$ is connected to exactly one vertex in H. This follows immediately from the fact that in the triangulation corresponding to v, the point p is skipped. It can be flipped to obtain a unique triangulation where p is part of some cell, whose corresponding vertex is therefore in H. We have shown that Lemma 5.13 implies that Theorem 5.11 holds for all configurations where $x + 1$ points are repeated.

We now prove Lemma 5.13. Using the Local Menger Lemma introduced by Wagner and Welzl [82], it is sufficient to show the existence of $k + 1 - \ell$ paths between pairs of vertices at distance 2.

Lemma 5.14 (Local Menger [82]). *Let $k \geq 2$ be an integer and let G be a connected simple undirected graph. The graph G is k-vertex connected if and only if G has at least $k+1$ vertices and for any pair of vertices u and v at distance 2 there are k pairwise internally vertex disjoint u-v-paths.*

Proof of Lemma 5.13. The proof is illustrated in Figure 5.6. We use the same notation as in the definition of (H, ℓ)-clone. By assumption, for any $m + 1 \leq j_1 \leq n$, there exists a unique $1 \leq j_2 \leq m$, such that for any $1 \leq i \leq \ell$, $\{v_{0,j_1}, v_{i,j_2}\}$ is in E'. We define by ϕ the function that maps such a j_1 to this uniquely defined j_2.

We denote by G' the (H, ℓ)-clone of G. By Menger's Theorem, the statement is equivalent to say that for any pair of distinct vertices in G', there are $k + \ell - 1$ internally vertex-disjoint paths between them. Here, we use the local version of this theorem (Lemma 5.14), and assume that the two vertices of the pair are at distance 2 in G'. We are going to distinguish cases depending on whether the vertices belong to V_0 or to V_i for $i > 0$. Note that we only use the fact that the vertices are at distance 2 in the last case. Let us begin with the case when they both belong to V_0. This is illustrated in Figure 5.6a.

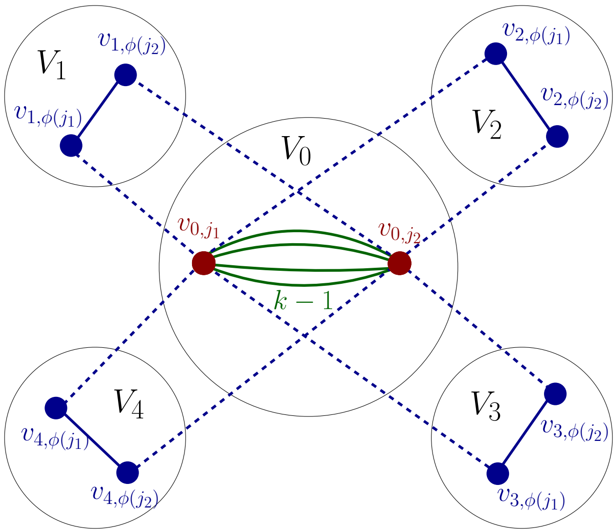

(a) Both vertices are in V_0

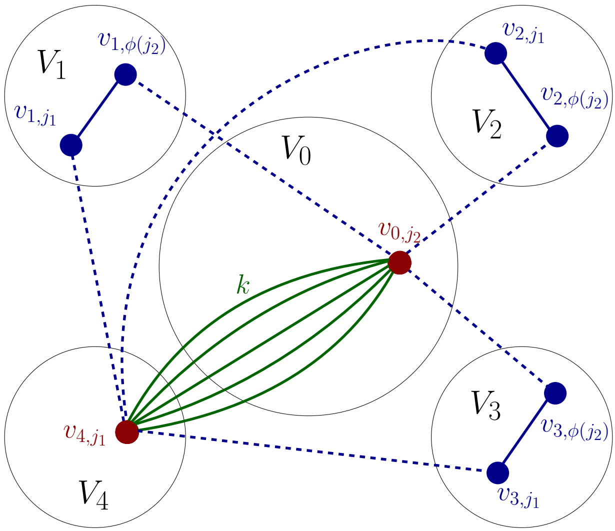

(b) One vertex is in V_0 and the other in V_4

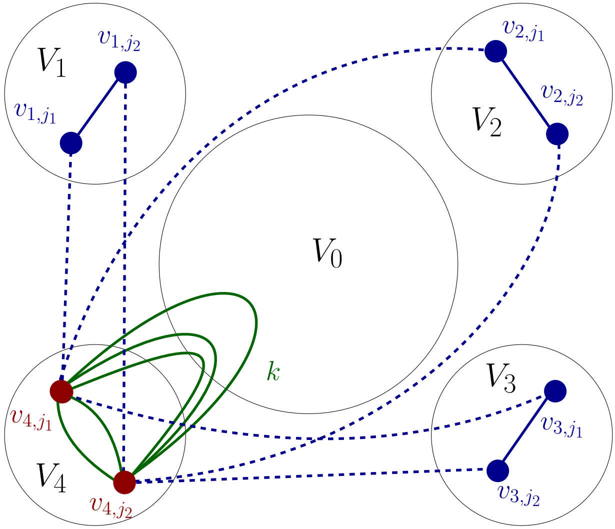

(c) Both vertices are in V_4

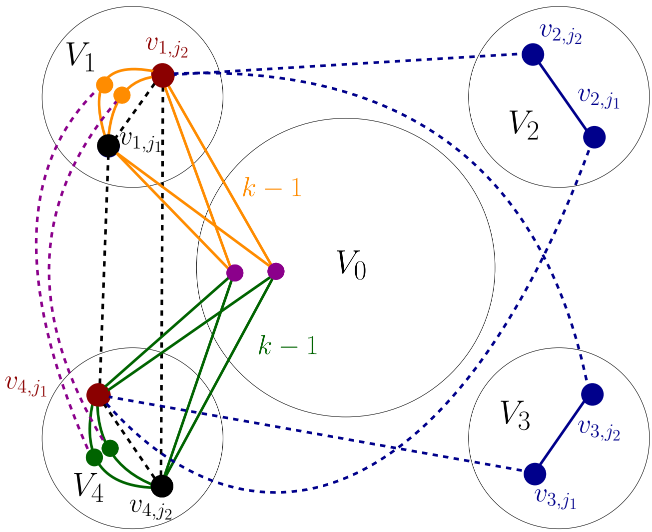

(d) Different vertices in H in different copies of H

Figure 5.6: Illustration of Lemma 5.13. Dashed lines represent edges. Thick lines represent paths of arbitrary lengths. Such a path can even be of length 0, as the endpoints might be the same vertex. The green and orange numbers denote the number of green and orange paths, respectively.

Let $m+1 \leq j_1, j_2 \leq n$ be two different integers. We show that there are $k+\ell-1$ (internally vertex-disjoint) paths between v_{0,j_1} and v_{0,j_2}. By assumption, the graph induced by V_0 is $(k-1)$-connected. Therefore there are $k-1$ paths between the two vertices for which all vertices are in V_0. Moreover by assumption, for all $1 \leq i \leq \ell$, $\{v_{0,j_1}, v_{i,\phi(j_1)}\}$ and $\{v_{0,j_2}, v_{i,\phi(j_2)}\}$ are in E'. Note that $v_{i,\phi(j_1)}$ and $v_{i,\phi(j_2)}$ might be the same. As H is connected, there is a path between $v_{i,\phi(j_1)}$ and $v_{i,\phi(j_2)}$ that uses only vertices of V_i. In total, we have found $(k-1)+\ell$ paths.

Let us now assume that one vertex belongs to V_0, and the other to V_{i_1} for some $1 \leq i_1 \leq \ell$. This is illustrated in Figure 5.6b. Let $1 \leq j_1 \leq m$ and $m+1 \leq j_2 \leq n$ be two integers. We show that there are $k+\ell-1$ paths between v_{i_1,j_1} and v_{0,j_2}. As G is k-connected, there are k paths between the two for which all vertices are in $V_0 \cup V_{i_1}$. Moreover, for all $1 \leq i_2 \leq \ell$ with $i_2 \neq i_1$, $\{v_{i_1,j_1}, v_{i_2,j_1}\}$ is in E'. Furthermore, $\{v_{0,j_2}, v_{i_2,\phi(j_2)}\}$ is in E'. As H is connected, there is a path between v_{i_2,j_1} and $v_{i_2,\phi(j_2)}$, which only uses vertices of V_{i_2}. There are $\ell-1$ such paths, which makes in total $k+(\ell-1)$ paths.

Now, we assume that the vertices belong to the same V_{i_1}, for $1 \leq i_1 \leq \ell$. This is illustrated in Figure 5.6c. Let $1 \leq j_1, j_2 \leq m$ be some different integers. We show that there are $k+\ell-1$ paths between v_{i_1,j_1} and v_{i_1,j_2}. As G is k-connected, there are k paths between them which only use vertices of $V_0 \cup V_{i_1}$. For all $1 \leq i_2 \leq \ell$, with $i_2 \neq i_1$, $\{v_{i_1,j_1}, v_{i_2,j_1}\}$ and $\{v_{i_1,j_2}, v_{i_2,j_2}\}$ are in E'. Moreover, as H is connected, there is a path between v_{i_2,j_1} and v_{i_2,j_2} that only uses vertices of V_{i_2}. In total, there are $k+(\ell-1)$ paths.

Finally, let us assume that the vertices belong to V_{i_1} and V_{i_2}, with $i_1 \neq i_2$. Let $1 \leq j_1, j_2 \leq m$ be some integers. We show that there are $k+\ell-1$ paths between v_{i_1,j_1} and v_{i_2,j_2}. Here, we use the assumption that v_{i_1,j_1} and v_{i_2,j_2} are at distance 2. This immediately implies $j_1 \neq j_2$, otherwise the two vertices would be adjacent. Let us consider a vertex v in V', that is adjacent to both v_{i_1,j_1} and v_{i_2,j_2}, and that belongs to V_{i_3} for some $0 \leq i_3 \leq \ell$. If $i_3 = 0$, then v is a vertex of H, so we use the fact that v is adjacent to a unique vertex in $G \setminus H$ to infer $j_1 = j_2$, a contradiction. Likewise, if $i_3 \notin \{0, i_1, i_2\}$, then by definition of a clone we have $j_1 = j_2$.

Thus, we have that each vertex v adjacent to both v_{i_1,j_1} and v_{i_2,j_2} is in V_{i_1} or in V_{i_2}. By assumption, there exists at least one such v, and we assume without loss of generality that v is in V_{i_1}. As there exists a unique vertex in V_{i_1} adjacent to v_{i_2,j_2}, we infer that v is equal to v_{i_1,j_2}. This is illustrated in Figure 5.6d. As G is k-connected, there are k paths between v_{i_1,j_1} and v_{i_1,j_2} that only use vertices of $V_0 \cup V_{i_1}$. Without loss of generality, we can assume that the path consisting of the single edge $\{v_{i_1,j_1}, v_{i_1,j_2}\}$ is one of them. We denote this path consisting of a single edge by P_e. For any path P_1 among the remaining $k-1$ paths, we can associate to each vertex of P_1 in V_{i_1}, a corresponding vertex in V_{i_2}. As the graphs induced by $V_0 \cup V_{i_1}$ and $V_0 \cup V_{i_2}$ are isomorphic, this characterises a path P_2 between v_{i_2,j_1} and v_{i_2,j_2} that uses only vertices of $V_0 \cup V_{i_2}$. We associate P_1 and P_2 in order to obtain a path from v_{i_1,j_1} to v_{i_2,j_2}. Observe that if there is a vertex v_0 of P_1 in V_0, then v_0 is also an internal vertex of P_2. Thus, we obtain a path by following P_1 until v_0, and then following P_2 to v_{i_2,j_2}. If there is no such vertex of P_1 in V_0, let us consider one internal vertex v_1 of P_1 in V_{i_1}. Such a vertex exists, because we assume P_1 not to be equal to P_e. Therefore, there is a vertex v_2 in V_{i_2}, that is an internal vertex of P_2, with $\{v_1, v_2\} \in E'$. We obtain a path by following P_1 until v_1, then going to v_2, and finally following P_2 to v_{i_2,j_2}. Observe that by doing so we have found $k-1$ internally vertex disjoint paths from v_{i_1,j_1} to v_{i_2,j_2} which only use vertices in $V_0 \cup V_{i_1} \cup V_{i_2} \setminus \{v_{i_1,j_2}, v_{i_2,j_1}\}$. We add the two following paths of length 2: $(v_{i_1,j_1}, v_{i_1,j_2}, v_{i_2,j_2})$ and $(v_{i_1,j_1}, v_{i_2,j_1}, v_{i_2,j_2})$. Finally, for each $1 \leq i_3 \leq \ell$, with $i_3 \neq i_1$ and $i_3 \neq i_2$, the edges $\{v_{i_1,j_1}, v_{i_3,j_1}\}$ and $\{v_{i_2,j_2}, v_{i_3,j_2}\}$ are in E'. As H is connected, there is a path between v_{i_3,j_1} and v_{i_3,j_2} that only uses vertices of V_{i_3}. Therefore we obtain $\ell - 2$ new paths. In total there are $(k-1)+2+(\ell-2)=k+\ell-1$ paths. □

5.3 Some properties of subdivisions in any dimension

The connectivity results we present in this paper hold in dimension 2. It is known that they do not hold in dimension 5 and higher [72, 73]. However, some of our lemmas hold in all dimensions. We state them in this section. Recall that we have shown in Section 5.2 why we can afford to only consider point configurations where no points are repeated, which can simply be seen as point sets. In particular, we make no difference between a label and its corresponding point. The following definitions are inspired by [82]. An important difference is that we allow for degenerate point sets, whereas in [82] only point sets in the plane in general position are considered. Let $\mathcal{P}$ be a set of n points in $\mathbb{R}^d$, such that the dimension of the affine hull of P is d. A point p in $\mathcal{P}$ is *extreme* if $\mathcal{P}$ and $\mathcal{P}\setminus\{p\}$ do not have the same convex hull. Observe that a point can lie on the boundary of the convex hull while still not being extreme.

Definition 5.15 (Bystanders)**.** Let $\mathcal{S}$ be a subdivision. If a label j is in some cell of $\mathcal{S}$, but $\{j\}$ is not a cell in $\mathcal{S}$, then j is called a *bystander.*

It follows easily from the definition of subdivisions that j is a bystander if and only if in all cells that contain j it is not an extreme point. Thus a point is either skipped, a bystander, or an extreme point in all cells that contain it. In the last case, we call it an *involved point.* Observe that skipping all bystanders in a subdivision $\mathcal{S}$ gives a refinement of $\mathcal{S}$. Recall that we assume here that no point is repeated. If we had allowed repeated points and kept the same definition, a repeated point would be a bystander. But if all points are repeated then skipping all of them would not yield a subdivision.

This is one of the reasons why we treated the case with repeated points in Section 5.2, and from now on forbid repeated points. We call *size* of a d-cell B the number of involved points in B.

Definition 5.16 (Slack of a subdivision)**.** Let $\mathcal{S}$ be a subdivision whose affine hull has dimension d. The *slack* of $\mathcal{S}$ is equal to $(\sum_i f_i - d - 1) + \beta$, where f_i denotes the number of d-cells of size i and β the number of bystanders.

Lemma 5.17. *A subdivision is an almost-triangulation if and only if it has slack* 1*.*

Proof. Observe that a subdivision is an almost-triangulation if and only if it has exactly one d-cell of size $d+2$ and the others of size $d+1$, and no bystander, or if all d-cells have size $d+1$ and there is one bystander. This is exactly the definition of being of slack 1. □

Lemma 5.18. *A subdivision has height* 2 *if and only if it has slack* 2*. In particular, any proper refinement of a subdivision of slack* 2 *has slack at most* 1*.*

Proof. Let $\mathcal{S}$ be a subdivision of height 2. Observe that it has at most two bystanders, and that all d-cells are of size at most $d+3$. If there are two bystanders, then all d-cells must be of size $d+1$. If there is exactly one bystander, then as $\mathcal{S}$ is not an almost-triangulation, there must be a d-cell of size $d+2$, and all others must be of size $d+1$. If there are no bystanders, then there is either a d-cell of size $d+3$ or two d-cells of size $d+2$. In all cases we have a subdivision of slack 2. □

By Lemmas 5.8 and 5.17, we know that a proper refinement of a subdivision of slack 1 has slack equal to 0. Lemma 5.18 could make us believe that this is a general property, i.e. that a proper refinement of a subdivision of slack k has slack at most $k-1$. However, this is not the case: See [82] for an example of a subdivision of slack 5 with a refinement of slack 6.

Lemma 5.19. *Let $\mathcal{S}$ and $\mathcal{S}'$ be two distinct almost-triangulations. If a subdivision $\mathcal{S}_2$ at height 2 coarsens $\mathcal{S}$ and $\mathcal{S}'$, then it is unique.*

Proof. In an almost-triangulation, there are two possibilities. Either there is a bystander and all d-faces have size $d+1$, or there is no bystander and all faces have size $d+1$ except for one that has size $d+2$. If p is a bystander in $\mathcal{S}$ and q is a bystander in $\mathcal{S}'$, then $\mathcal{S}_2$ is the unique subdivision where p and q are bystanders and all d-faces have size $d+1$. Likewise, if B is a d-face of size $d+2$ in $\mathcal{S}$, then $\mathcal{S}_2$ is the unique subdivision where q is a bystander and all d-faces have size $d+1$ except for B that has size $d+2$. Finally, let us assume that B' is a d-face of size $d+2$ in $\mathcal{S}'$. If $B \cap B'$ contains at most d points, then $\mathcal{S}_2$ is the unique subdivision where all d-faces have size $d+1$ except for B and B' of size $d+2$. Lastly, if $B \cap B'$ contains $d+1$ points (it cannot contain $d+2$ as $\mathcal{S} \neq \mathcal{S}'$), then $\mathcal{S}_2$ is the unique subdivision where all d-faces have size $d+1$ except for $B \cup B'$ which has size $d+3$. □

Definition 5.20 (Set of refining triangulations)**.** For a subdivision $\mathcal{S}$, we denote by $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$ the set of triangulations that refine $\mathcal{S}$.

For a triangulation $\mathcal{T}$ and an almost-triangulation $\mathcal{S}$ that coarsens $\mathcal{T}$, we denote by $\mathcal{T}[\mathcal{S}]$ the unique triangulation distinct from $\mathcal{T}$ that refines $\mathcal{S}$. Its existence and uniqueness is given by Lemma 5.8. Observe that if $\mathcal{S}_1$ and $\mathcal{S}_2$ are different almost-triangulations that coarsen $\mathcal{T}$, then $\mathcal{T}[\mathcal{S}_1]$ and $\mathcal{T}[\mathcal{S}_2]$ are different triangulations.

Lemma 5.21. *i) For two subdivision $\mathcal{S}_1$ and $\mathcal{S}_2$ of slack 2, $\mathbf{T}_{ref}\langle\mathcal{S}_1\rangle \cap \mathbf{T}_{ref}\langle\mathcal{S}_2\rangle$ is either a) empty, b) equals $\{\mathcal{T}\}$ for some triangulation $\mathcal{T}$, c) equals $\mathbf{T}_{ref}\langle\mathcal{S}_3\rangle$ for some almost-triangulation $\mathcal{S}_3$, or d) $\mathcal{S}_1 = \mathcal{S}_2$.*

ii) Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two different almost-triangulations that coarsen *a triangulation* $\mathcal{T}$. *If there is a subdivision* $\mathcal{S}_3$ *of slack* 2 *that satisfies* $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_2]\} \subseteq \mathbf{T}_{ref}\langle\mathcal{S}_3\rangle$, *then this subdivision* $\mathcal{S}_3$ *is unique.*

Proof. i) Let us assume that $\mathbf{T}_{ref}\langle\mathcal{S}_1\rangle \cap \mathbf{T}_{ref}\langle\mathcal{S}_2\rangle$ is not empty. By Lemma 5.18, $\mathcal{S}_1$ and $\mathcal{S}_2$ have height 2. Let us denote by $\mathcal{S}$ a refinement of $\mathcal{S}_1$ and $\mathcal{S}_2$ with maximal height. If $\mathcal{S}$ has height 0 we are in case b). If it has height 2 we are in case d). Let us now assume that it has height 1. By Lemma 5.19, if another subdivision $\mathcal{S}'$ refines $\mathcal{S}_1$ and $\mathcal{S}_2$, then $\mathcal{S}'$ is equal to $\mathcal{S}$. Therefore we are in case c).

ii) Let $\mathcal{S}_3$ and $\mathcal{S}_4$ be two subdivisions of slack 2 such that we have $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_2]\} \subseteq \mathbf{T}_{ref}\langle\mathcal{S}_3\rangle$ and $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_2]\} \subseteq \mathbf{T}_{ref}\langle\mathcal{S}_4\rangle$. It implies that we have $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_2]\} \subseteq \mathbf{T}_{ref}\langle\mathcal{S}_3\rangle \cap \mathbf{T}_{ref}\langle\mathcal{S}_4\rangle$. Since we are not in case a), b) or c), we have $\mathcal{S}_3 = \mathcal{S}_4$. □

Definition 5.22 (Compatible almost-triangulations)**.** Let $\mathcal{T}$ be a triangulation. Two almost-triangulations $\mathcal{S}_1$ and $\mathcal{S}_2$ that coarsen $\mathcal{T}$ are called *compatible with respect to* $\mathcal{T}$, in symbols $\mathcal{S}_1 \diamond \mathcal{S}_2$ if there exists a subdivision $\mathcal{S}_{1,2}$ of slack 2 such that $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_2]\} \subseteq \mathbf{T}_{ref}\langle\mathcal{S}_{1,2}\rangle$. Otherwise, $\mathcal{S}_1$ and $\mathcal{S}_2$ are called *incompatible with respect to* $\mathcal{T}$, in symbols $\mathcal{S}_1 \not\diamond\, \mathcal{S}_2$

Observe that by Lemma 5.21 that if $\mathcal{S}_{1,2}$ exists then it is unique. We now observe that for any triangulation $\mathcal{T}$, and any two almost-triangulations $\mathcal{S}_1$ and $\mathcal{S}_2$ that coarsen $\mathcal{T}$ with $\mathcal{S}_1 \diamond \mathcal{S}_2$, there exists a $\mathcal{T}$-avoiding $\mathcal{T}[\mathcal{S}_1]$-$\mathcal{T}[\mathcal{S}_2]$-path of length at most $d+1$ in the bistellar flip graph. It is sufficient to show the following lemma:

Lemma 5.23. *Let* $\mathcal{S}$ *be a subdivision of slack* 2. *Then the subgraph of the bistellar flip graph induced by* $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$ *is a* k*-cycle, where* $3 \leq k \leq d+3$.

Proof. First, it is easy to observe that if there are exactly two d-cells of size $d+1$ that contain a bystander, or if there are two d-cells of size $d+2$, or if there is one d-cell B of size $d+2$ and one bystander

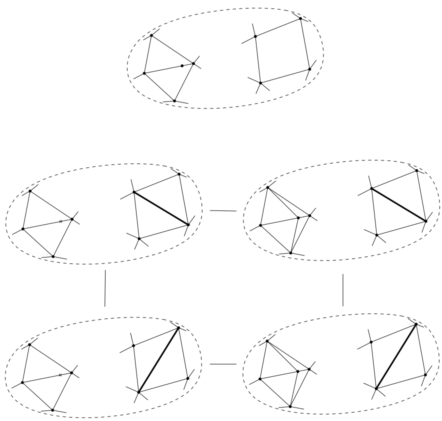

Figure 5.7: A 2-dimensional subdivision $\mathcal{S}$ of slack 2, and the four triangulations in $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$. The flips can be done independently, and so $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$ induces a 4-cycle.

not in B, then there are four triangulations in $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$ and they form a cycle of length 4, as shown in Figure 5.7 for the case $d = 2$. If we are not in one of the previous cases, there is a d-face B of size $d + 3$. Observe that we can focus on B and forget about the rest of the subdivision. Chapter 5.5.1 in [22] deals with configurations of $d + 3$ points. It is shown that all subdivisions of a d-dimensional point configuration with $d+3$ points are regular. Thus, the bistellar flip graph of this point configuration is the graph of a $(n - d - 1)$-dimensional polytope, that is a polygon. Moreover, it is shown that there are at most $d+3$ triangulations. For the sake of completeness, we illustrate in Figure 5.8 the situation for $d = 2$. The cases where no three points are on a line were already presented in [82]. □

Recall that by Lemma 5.8 for a triangulation $\mathcal{T}$ and an almost-triangulation $\mathcal{S}$ that coarsens $\mathcal{T}$, there exists a corresponding triangulation $\mathcal{T}'$ that refines $\mathcal{S}$, such that one can flip from $\mathcal{T}$ to $\mathcal{T}'$. We now claim that there is an equivalent thing between subdivisions of

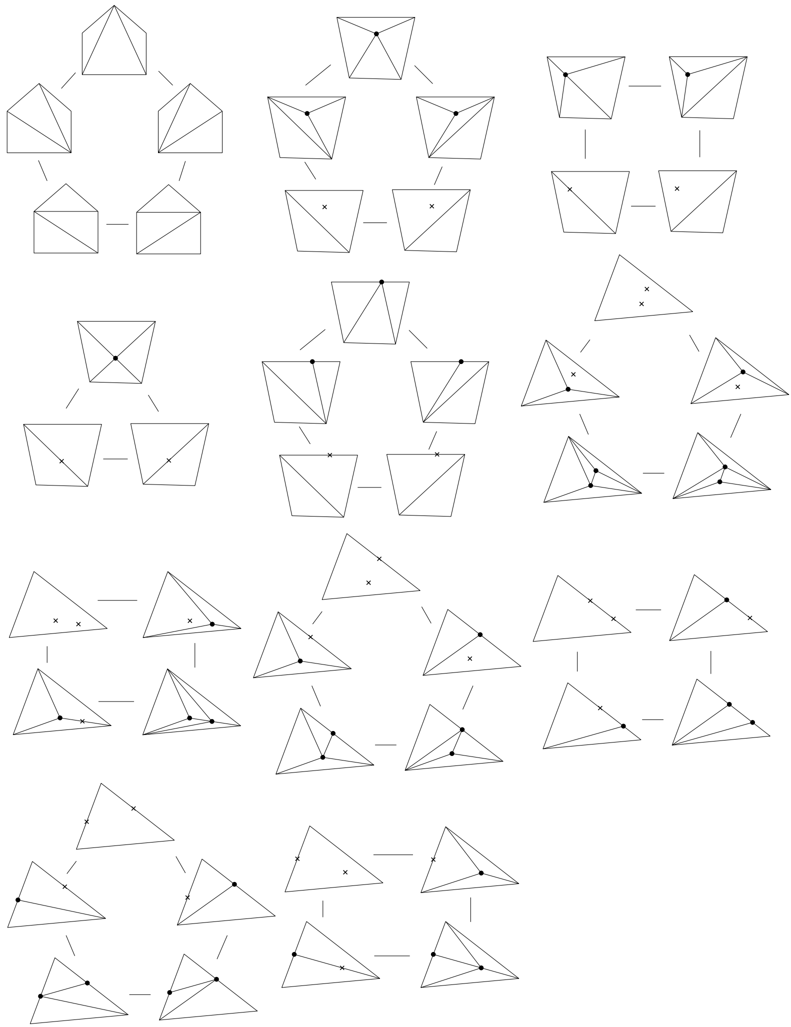

Figure 5.8: Bistellar flip graphs for five points in the plane. A cross indicates a skipped point.

slack 1 and 2.

Lemma 5.24. *Let $\mathcal{T}$ be a triangulation, and let $\mathcal{S}_1$ be an almost-triangulation that coarsens $\mathcal{T}$. For any subdivision $\mathcal{S}_2$ of slack 2 that coarsens $\mathcal{S}_1$, there exists an almost-triangulation $\mathcal{S}_1'$ that coarsens $\mathcal{T}$ and refines $\mathcal{S}_2$ such that $\mathcal{S}_1 \diamond \mathcal{S}_1'$ with respect to $\mathcal{T}$. Also, distinct coarsenings of $\mathcal{S}_1$ lead to distinct almost-triangulations.*

Proof. By Lemma 5.23, $\mathbf{T}_{ref}\langle \mathcal{S}_2 \rangle$ is a k-cycle, where $3 \leq k \leq d+3$. There are two triangulations which are adjacent to $\mathcal{T}$ in this cycle. As $\mathcal{T}[\mathcal{S}_1]$ is one, the other is uniquely defined, and can be written as $\mathcal{T}[\mathcal{S}_1']$ for some almost-triangulation $\mathcal{S}_1'$. By Lemma 5.21, if $\{\mathcal{T}[\mathcal{S}_1], \mathcal{T}, \mathcal{T}[\mathcal{S}_1']\} \subseteq \mathbf{T}_{ref}\langle \mathcal{S}_2' \rangle$ for some subdivision of slack 2 $\mathcal{S}_2'$, then we have $\mathcal{S}_2' = \mathcal{S}_2$. This implies that distinct coarsenings of $\mathcal{S}_1$ lead to distinct almost-triangulations. □

Lemma 5.25. *Let $\mathcal{T}$ be a triangulation, and let p be a skipped point in $\mathcal{T}$. Let $\mathcal{S}_1$ denote the almost-triangulation identical to $\mathcal{T}$ except that p is now a bystander. Let $\mathcal{S}_2$ be any almost-triangulation that coarsens $\mathcal{T}$ and which is not $\mathcal{S}_1$. We have that $\mathcal{S}_1$ and $\mathcal{S}_2$ are compatible.*

Proof. Observe that p must be skipped in $\mathcal{S}_2$, for otherwise it would be equal to $\mathcal{S}_1$. Now let us consider $\mathcal{S}_3$ which is identical to $\mathcal{S}_2$ except that p is now a bystander. Therefore $\mathcal{S}_3$ has slack 2. By assumption, $\mathcal{S}_3$ coarsens $\mathcal{T}$ and $\mathcal{T}[\mathcal{S}_2]$. Now, for the same reason that $\mathcal{S}_2$ coarsens $\mathcal{T}$, we have that $\mathcal{S}_3$ coarsens $\mathcal{S}_1$. This is because the only modification is that p is a bystander in $\mathcal{S}_1$ and $\mathcal{S}_3$ whereas it is skipped in $\mathcal{T}$ and $\mathcal{S}_2$. Thus we have $\mathcal{S}_1 \diamond \mathcal{S}_2$. □

5.4 Definitions for point sets in the plane

In this section, we consider point sets in the plane. For some of our purposes, it will be convenient to look at a subdivision as straight-line plane graph. This is only possible because we are not allowing point repetitions. Thus there is no risk of confusion as to which point is an edge incident.

Definition 5.26 (The graph of a subdivision)**.** Let $\mathcal{S}$ be a subdivision. Let $\mathsf{V}^{\mathsf{inv}}$ be the set of involved points, and let V be the set of points that are not skipped. The graph $G = (\mathsf{V}, E)$ of $\mathcal{S}$ is the straight-line plane graph with vertex set V, where there is an edge between two vertices in $\mathsf{V}^{\mathsf{inv}}$ if they are the two extreme points of some 1-cell in $\mathcal{S}$.

Observe that the bystanders may lie on some edge, but they are still of degree 0. Furthermore, G is a convex partition on $\mathsf{V}^{\mathsf{inv}}$ (if we forget about the bystanders). Let us recall that the converse is not true, as shown in Figure 5.3: There are convex partitions which are not the plane graph of any subdivision. The graph of a triangulation is a maximal plane graph on $\mathsf{V} = \mathsf{V}^{\mathsf{inv}}$.

5.4.1 The unoriented edge Lemma

We now consider a fixed subdivision $\mathcal{S}$ on $\mathcal{P}$ with N involved points and no bystanders. We consider $G = (\mathsf{V}^{\mathsf{inv}}, E)$ its corresponding plane graph. Throughout Section 5.4.1, we say *edge* for an edge of G. The aim of this section is to show a lower bound on the number of $(2, 2)$-flips that can be done.

Definition 5.27 (Locked edges)**.** Let p be a point in $\mathcal{P}$, and e be an edge incident to p, with e not on the convex hull. We say that e is *concave-locked* at p if its removal would create an angle larger than π at p. Similarly, e is *straight-locked* at p if its removal would create an angle of exactly π. An edge is *locked* at p if it is concave-locked or straight-locked at p. An edge is *unlocked* if it is not locked at any endpoint.

Definition 5.28. We denote by $\mathcal{P}_{i,j}$ the set of inner points p in $\mathcal{P}$ with i edges concave-locked at p and j edges straight-locked at p. We denote by $C_{i,j}$ the cardinality of $\mathcal{P}_{i,j}$. Likewise, $\mathcal{P}^b_{i,j}$ denotes the set of points p on the convex hull with i edges concave-locked at p and j edges straight-locked at p, and $C^b_{i,j}$ denotes the cardinality of $\mathcal{P}^b_{i,j}$.

Observation 5.29. *If p is in $\mathcal{P}_{3,0}$, then there are exactly three edges incident to p. If p is in $\mathcal{P}_{1,2}$, there are exactly four edges incident to p, two of them being aligned. If p is in $\mathcal{P}_{0,4}$, there are also four edges incident to p, the angle between any two consecutive edges being $\pi/2$ (there are two pairs of aligned edges). If p is in $\mathcal{P}^b_{0,1}$, there are three edges incident to p, two of them being aligned and on the boundary of the convex hull.*

Definition 5.30 (Orientation of a subdivision)**.** Let $\mathcal{S}$ be a subdivision. An *orientation* of $\mathcal{S}$ is the assignment to each inner edge of either having no direction, or being directed towards one of its endpoints, or being directed towards both endpoints. We call these edges *unoriented*, *directed*, *mutual*, respectively. We denote by C_i the number of inner points with indegree i, and by C^b_i be the number of non-extreme points on the boundary of the convex hull with indegree i.

Lemma 5.31 (Unoriented edges Lemma)**.** *Let $\mathcal{S}$ be a subdivision on N involved points, no bystanders, and with D edges missing towards being a triangulation. We consider an orientation of $\mathcal{S}$. The number of unoriented inner edges is at least $N - 3 - D + \sum_i((2-i)C_i) - \sum_i iC^b_i$. The inequality becomes an equality under the assumption that edges can be oriented towards at most one endpoint.*

Proof. Let O denote the number of oriented inner edges, and U the number of unoriented inner edges. Let us denote by n° the number of inner points in $\mathcal{S}$, and by h the number of points on the boundary of the convex hull. There are $3N-3-h$ edges in a triangulation of the involved points. The number of inner edges in $\mathcal{S}$ is therefore equal to $3N-3-2h-D = N+2n^{\circ}-3-D$. We have $O \leq \sum_i i(C_i+C_i^b)$, and there is equality if no edge is oriented towards both endpoints. Therefore we obtain $U = N+2n^{\circ}-3-D-O \geq N-3-D+2n^{\circ} - \sum_i i(C_i+C_i^b) = N-3-D+\sum_i((2-i)C_i) - \sum_i iC_i^b$. □

Definition 5.32. We say that a partial orientation on the inner edges of a subdivision $\mathcal{S}$ is *admissible* if we have the following properties:

1. $C_i^b > 0 \implies i \leq 1$,
2. $C_i > 0 \implies i \leq 4$.

Observation 5.33. *Let $\mathcal{S}$ be a subdivision. Orient all inner edges towards points where they are locked. This orientation is admissible.*

Lemma 5.34. *In an admissible orientation of a subdivision, there are at least $N-3-D+2C_0+C_1-2C_4-C_3-C_1^b$ unoriented inner edges.*

Proof. This is a direct application of Lemma 5.31 and the definition of being admissible. Let us denote by U the number of unoriented inner edges. We have $U \geq N-3-D+\sum_i((2-i)C_i) - \sum_i iC_i^b = N-3-D+2C_0+C_1-C_3-2C_4-C_1^b$. □

5.4.2 Coarseners of subdivisions

We now define the coarseners of a subdivision, similarly to [82]. When no three points are on a line, there are only two kinds of flips: $(2,2)$-flips and $(3,1)$-flips. A $(2,2)$-flip is characterised by giving the edge which is to be flipped. A $(3,1)$-flip is characterised by giving which point is to be flipped. This is not true for $(2,1)$-flips, which can happen when several points are allowed to be on a line. When making a point a bystander, one must also specify which edges are to be removed. This issue also appears when defining coarseners. Wagner and Welzl characterise a coarsener simply by giving a set of points. Below, we extend their definition to our broader setting, where we need to specify a set of edges.

Definition 5.35 (Coarsener)**.** *Let $\mathcal{S}$ be a subdivision. Let U be a non-empty set of non-extreme points, containing only involved points, and let E be a set of* 1*-cells that each contain at least one point in U. We say that (U, E) is a* coarsener *if there exists a coarsening $\mathcal{S}'$ of $\mathcal{S}$ such that:*

1. *Every point in U is a bystander in $\mathcal{S}'$, and every involved point not in U remains involved in $\mathcal{S}'$.*
2. *Every* 1*-cell in E is not a subset of any* 1*-cell in $\mathcal{S}'$,*
3. *Every* 1*-cell not in E which contains a point in U is a proper subset of a* 1*-cell in $\mathcal{S}'$.*
4. *Every* 1*-cell not in E which does not contain a point in U is a* 1*-cell of $\mathcal{S}'$.*

We say that $\mathcal{S}'$ is the subdivision obtained by *isolating the coarsener* (U, E). Every subdivision with at least one non-extreme point involved has a coarsener. Simply take U as the set of all non extreme points, and E as the set of all inner edges. By isolating this coarsener, we obtain the trivial subdivision $\mathsf{S}_{\mathsf{triv}}$.

Let us analyse the implications of Definition 5.35. A point u in U must become a bystander in $\mathcal{S}'$ by Property 1., therefore u may be either part of a 1-cell or a 2-cell. If u lies on the boundary of the convex hull, then it must be part of a 1-cell. By Property 2. and the fact that $\mathcal{S}'$ is a coarsening of $\mathcal{S}$, a 1-cell in E is a proper subset of a 2-cell in $\mathcal{S}'$. Visually, the corresponding edge in the straight-line plane graph has disappeared. Relatedly, the two aligned edges corresponding to 1-cells not in E which contain u are now merged into one edge on which the point u lies. This is just a visual description, which is not formal. By Property 4. all the other edges remain the same. Assuming that E contains all edges needed such that every point in U becomes isolated, either lying on a 2-cell or on a 1-cell, the only reason why $\mathcal{S}'$ may not exist would be because some 2-cells are not convex anymore.

Definition 5.36 (Fixing a pair of 1-cells)**.** Let $\mathcal{S}$ be a subdivision, and let (U, E) be a coarsener with corresponding coarsening $\mathcal{S}'$. Let u be a point in U that is part of some 1-cell in $\mathcal{S}'$. The two 1-cells in $\mathcal{S}$ not in E that contain u are called *fixed by the coarsener* (U, E) *at* u. For a point u that is not contained in any 1-cell in $\mathcal{S}'$, we say that (U, E) *fixes no* 1*-cell at* u.

Recall that if two 1-cells are fixed by a coarsener at point u, then all the points they contain are aligned. Therefore, if two coarseners fix the same 1-cell at u, then they fix the same pair of 1-cells at u.

Definition 5.37 (Increment of a coarsener)**.** Let $\mathcal{S}$ be a subdivision, and (U, E) be a coarsener with corresponding coarsening $\mathcal{S}'$. Let $U_1 \subseteq U$ denote the set of inner points in U that are contained in some 1-cell in $\mathcal{S}'$. Let $U_1^b \subseteq U$ the set of points, non-extreme but lying on the boundary of the convex hull, that are contained in some 1-cell in $\mathcal{S}'$. The *increment* of (U, E), denoted by $\mathsf{inc}(U, E)$, is equal to $|E| - 2|U| + |U_1| + 2|U_1^b|$.

Lemma 5.38. *Let $\mathcal{S} = (V, E)$ be a subdivision, and (U, E) be a coarsener with corresponding coarsening $\mathcal{S}'$. We have $\mathsf{sl}\mathcal{S}' = \mathsf{sl}\mathcal{S} + \mathsf{inc}(U, E)$.*

Proof. By Definition 5.16, $\mathsf{sl}\mathcal{S}$ is equal to $(\sum_i f_i - 3) + \beta$, where f_i denotes the number of 2-cells of size i and β the number of bystanders. Remark that $\sum_i f_i - 3$ corresponds to the number of 1-cells missing towards being a triangulation. Let us denote by N the number of points involved in $\mathcal{S}$, and by h the number of points involved that lie on the boundary of the convex hull. Let us denote by $Y_{\mathcal{S}}$ the number of 1-cells in $\mathcal{S}$. Therefore we have $\mathsf{sl}\mathcal{S} = 3N - 3 - h - Y_{\mathcal{S}} + B$. Each point in U is a bystander in $\mathcal{S}'$. Thus the number of bystanders in $\mathcal{S}'$ is $B' := B + |U|$, and the number of involved points is $N' := N - |U|$. For the same reason, the number of involved points on the boundary of the convex hull in $\mathcal{S}'$ is $h' := h - |U_1^b|$. Each 1-cell in E does not appear in $\mathcal{S}'$, and for each point in $U_1 \cup U_1^b$, two 1-cells in $\mathcal{S}$ correspond to only one 1-cell in $\mathcal{S}'$. Thus the number of 1-cells in $\mathcal{S}'$ is $Y'_{\mathcal{S}} = Y_{\mathcal{S}} - |E| - |U_1| - |U_1^b|$. Finally, we have $\mathsf{sl}\mathcal{S}' = 3N' - 3 - h' - Y'_{\mathcal{S}} + B' = (3N - 3|U|) - 3 - (h - |U_1^b|) - (Y_{\mathcal{S}} - |E| - |U_1| - |U_1^b|) + (B + |U|) = \mathsf{sl}\mathcal{S} + |E| - 2|U| + |U_1| + 2|U_1^b|$. □

Definition 5.39 (Prime coarsener)**.** Let $\mathcal{S}$ be a subdivision, and (U, E) be a coarsener. We say that (U, E) is a *prime* coarsener if for any set (U', E') with $E' \subseteq E$, we have that (U', E') is a coarsener if and only if $U' = U$ and $E' = E$.

Observation 5.40. Taking the notation of the Definition 5.39, we have that if (U', E') is a coarsener and $E' \subseteq E$ then $U' \subseteq U$. Indeed it is not possible to obtain a new bystander with fewer edges. Therefore the definition essentially means that there is no smaller coarsener.

Definition 5.41 (Perfect coarsener)**.** A prime coarsener (U, E) in a subdivision $\mathcal{S}$ is a *perfect coarsener* if $\mathsf{inc}(U, E) = 1$.

Definition 5.42 (Perfect coarsening)**.** A coarsening of a subdivision $\mathcal{S}$ is a *perfect coarsening* if it is obtained by either adding a skipped point as a bystander, removing an unlocked edge or by isolating a perfect coarsener.

Lemma 5.43. *Let $\mathcal{S}$ be a subdivision. Let $X_1 = (U_1, E_1)$ and $X_2 = (U_2, E_2)$ be coarseners. Let U be the set of points u in $U_1 \cap U_2$ such that either X_1 and X_2 fix the same pair of 1-cells at u, or at least one of them fixes no 1-cell. Let E be the set of 1-cells in $E_1 \cap E_2$ that contain at least one point in U. If U is not empty, then (U, E) is a coarsener.*

Proof. By assumption, there is at least one involved point u in U. For each such point, we have to show that either each 1-cell that contains u is in E, or that exactly two of such 1-cells B and B' are not in E, and then the points in B and B' are aligned. This implies that u is indeed a bystander in the coarsening $\mathcal{S}'$, obtained by isolating (U, E). If X_1 and X_2 fix the same pair of 1-cells at u, all other 1-cells that contain u are in E_1 and in E_2, and thus also in E. If one coarsener, say X_1, fixes no 1-cell and the other, X_2, fixes a pair B and B' of 1-cells, then every 1-cell that contains u apart from B and B' is in E. In the last case, both X_1 and X_2 fix no 1-cells at u, and then all 1-cells that contain u are in E.

It remains to show that the plane graph that we obtain from the new subdivision is a convex partition. Let u be a point in $\mathcal{P} \setminus (U_1 \cap U_2)$. Assume without loss of generality that u is not in U_1. We know that (U_1, E_1) is a coarsener. By isolating it, we obtain $\mathcal{S}_1$. As u is not in U_1, it is not a bystander in $\mathcal{S}_1$, and therefore the angle between any two consecutive 1-cells in $\mathcal{S}_1$ around u is at most π. There are fewer 1-cells in E than in E_1, thus the angle between any two consecutive 1-cells in $\mathcal{S} \setminus E$ around u is at most π. Finally, let u be a point in $U_1 \cap U_2$, such that X_1 and X_2 do not fix the same pair of 1-cells at u. Let B_1 and B_1' be the two 1-cells containing u whose points are aligned that are not in E_1. Likewise, let B_2, B_2' be the two 1-cells containing u whose points are aligned that are not in E_2. Observe that the four distinct 1-cells B_1, B_1', B_2, B_2' are not in E. As they consist of two pairs of aligned edges, the angle between any two consecutive 1-cells in $\mathcal{S} \setminus E$ around u is less than π. We have shown that if U is not empty, then (U, E) is a coarsener. □

Lemma 5.44. *Let $\mathcal{S}$ be a subdivision. If $X_1 = (U_1, E_1)$ and $X_2 = (U_2, E_2)$ are prime coarseners, then either $X_1 = X_2$, or $\forall u \in U_1 \cap U_2$, X_1 and X_2 do not fix the same pair of 1-cells at u (which includes the case $U_1 \cap U_2 = \emptyset$).*

Proof. We define U and E as in the statement of Lemma 5.43. Let us assume $X_1 \neq X_2$, which implies that (U, E) is not equal to X_1. As X_1 is prime, Lemma 5.43 implies that U is empty. This means that for any point u in $U_1 \cap U_2$, X_1 and X_2 do not fix the same pair of edges at u. □

Observation 5.45. *Let $\mathcal{S} = (V, E)$ be a subdivision, and let G be its corresponding plane graph. If (U, E) is a prime coarsener, then the subgraph of G with vertex set U and edge set E is connected.*

5.5 The coarsening Lemma

In this section, we give a lower bound on the number of perfect coarsenings in a subdivision $\mathcal{S}$ in the plane. We denote by $G = (\mathsf{V}, E)$ the plane graph corresponding to $\mathcal{S}$. We make no difference between an edge of G and its corresponding 1-cell in $\mathcal{S}$. We denote by n the number of points in the point set (including the skipped points), and by N the number of points that are involved or bystanders in $\mathcal{S}$. We extend the coarsening Lemma by Wagner and Welzl to point sets in degenerate position. This Lemma gives a lower bound on the number of perfect coarsenings of $\mathcal{S}$.

In the proof of the coarsening Lemma, we use an algorithm that allows us to identify perfect coarseners, and to give an admissible orientation. We call this algorithm the *orienting algorithm.* Recall that by Lemma 5.34, there are at least $N-3-D+2C_0+C_1-2C_4-C_3-C_1^b$ unoriented edges. We show that for each of this unoriented edge, there is one perfect coarsening obtained by removing this edge. Furthermore, for each inner point of indegree 3 there is one corresponding perfect coarsener, for each inner point of indegree 4 there are two corresponding perfect coarseners, and for each point on the boundary with indegree 1, there is one corresponding perfect coarsener. We also show that all these coarsenings are distinct. In total, we obtain $N-3-D+2C_0+C_1$ perfect coarsenings. We conclude using the following observation:

Observation 5.46. *Any point that is skipped can always be inserted. Therefore there are $n-N$ perfect coarsenings of $\mathcal{S}$ consisting in adding a skipped point as a bystander.*

5.5.1 The orienting algorithm

Let $\mathcal{S}$ be a subdivision with no bystander. The orienting algorithm has three phases. In the first one, we orient every edge towards a point where it is locked. It may be that edges are locked at both endpoints, and are therefore mutual. Observe that this orientation is admissible.

The candidate components and the second phase

Let U be a set of points and E a set of inner edges, such that $U \neq \emptyset$ and each edge in E contains a least one point in U. Recall that we speak equivalently of edges and of 1-cells. Let U_0 denote the set of points u in U such that all 1-cells in $\mathcal{S}$ that contain u are in E. Let $U_1 \subseteq U$ denote the set of inner points u in U such that exactly two 1-cells B and B' that contain u are not in E, and the points in $B \cup B'$ are aligned. Let $U_1^b \subseteq U$ the set of points, non-extreme but lying on the boundary of the convex hull, such that exactly two 1-cells B and B' that contain u are not in E, and the points in $B \cup B'$ are aligned. We say that (U, E) is a *candidate component* if:

1. We have $U = U_0 \cup U_1 \cup U_1^b$,
2. An edge in E with exactly one endpoint in U is oriented towards it and is not mutual,
3. No edge in E is unoriented,
4. Let $E' \subseteq E$ be the set of mutual edges in E. The subgraph (U, E') is a spanning tree of U,
5. There are exactly three edges in E oriented towards a point in U_0, exactly two edges in E oriented towards a point in U_1, and exactly one edge in E oriented towards a point in U_1^b.

If (U, E) is a candidate component, then $|E| = 3|U_0| + 2|U_1| + |U_1^b| - (|U|-1) = 2|U_0| + |U_1| + 1 = 2|U| - |U_1| - 2|U_1^b| + 1$. This immediately implies that if (U, E) is a coarsener, then $\mathsf{inc}(U, E) = 1$. Moreover in a coarsener, after the first phase of the orienting algorithm, either no edge or all edges of a candidate component are in the edge set of the coarsener. This is because the mutual edges form a spanning tree of U, and if two points connected by a mutual edge of a candidate component must be both bystanders or both involved in a coarsening of $\mathcal{S}$. This implies that if a candidate component (U, E) is a coarsener, then it is a prime coarsener, and even a perfect coarsener as $\mathsf{inc}(U, E) = 1$.

Now let us describe the second phase of the orienting algorithm. We are going to modify the orientation of the edges, while maintaining the property that if a candidate component is a coarsener, then it is a perfect coarsener. At the end, the aim is that (U, E) is a candidate component if and only if it is a perfect coarsener. Thus the candidate components are going to grow (in order to become closer to being a coarsener), or disappear (if no perfect coarsener contains them).

Let (U, E) be a candidate component. Suppose q is a point not in U such that removing the edges in E creates an angle of at least π at q. We orient one of the edges e in E incident to q towards

q, thus making e mutual. The other edges in E incident to q are called the *witnesses of the extra new orientation* of e from p to q. We iterate as long as possible. This must end, as at each step we make one directed edge become mutual. The candidate components may grow, or disappear if the new vertex does not have the right indegree, or is adjacent to an unoriented edge that needs to be in E. When we are done, it is the end of the second phase.

Let us make some remarks. When removing all witnesses incident to a point p, the edges that remain are locked at p. In particular, the indegree of p is at most 4, and if it is, then there are four edges incident to p and the angle between to consecutive edge around p is $\pi/2$. In this case we say that p is in the *regular situation*. If p is an inner point not in the regular situation, the indegree of p is at most 3. If p is on the boundary of the convex hull, then its indegree is at most 1. Observe that if an unoriented edge e is incident to a point p, then p the indegree of p is at most 3 and it is 3 if and only if e is aligned with another edge incident to p.

Observe that we maintained the following property: If two points in a candidate component (U, E) are connected by a mutual edge in E, then either both or non are bystanders in a coarsening of $\mathcal{S}$. As in a candidate component the mutual edges form a spanning tree of U, we have that if a candidate component is a coarsener, then it is a prime coarsener. Observe a point u is in U_1 may be adjacent to a mutual edge e, as long as e is not in E.

The third phase

It is clear that from an edge e, we can find all the points that need to be bystanders in a coarsening where e is missing. One has simply to do a depth-first search on the mutual edges from the endpoints of e, by applying the following rule. Let us assume that we are following a mutual edge from u to v. If when removing u and the witnesses of the extra new orientation of $\{u, v\}$ from u to v (if there are some) we obtain an angle larger than π at v, then we continue the search on all the mutual edges incident to v. This is because in any coarsening where u is a bystander, v must be a bystander lying in some 2-cell, and cannot be in any 1-cell. Now let us assume that we obtain an angle of exactly π, formed by the aligned edges denoted by e and e'. Observe that there cannot be mutual edges apart from $\{u, v\}$ on the same side of the edges e, e' as $\{u, v\}$. On the other side, there can be at most one mutual edge, denoted by e'', and let us denote by w the other endpoint of e''. Let us remove e'' and the witnesses of the extra new orientation of e from w to v (if there are some). If the angle at v is larger than π, then we continue the search on all the mutual edges incident to v. This is because in a coarsening where u is a bystander, v has to be a bystander, and thus so does w. But now, since the angle when removing w and its corresponding witnesses is larger than π, v cannot be a bystander in some 1-cell in the coarsening. Likewise, if when removing w and the witnesses of e'' we obtain an angle of π, but it is formed by a pair of aligned edges which is not e, e', then we continue the search on all the mutual edges incident to v. Once again, the reason is that v cannot be a bystander belonging to some 1-cell in the coarsening. Finally, if the pair of aligned edges is e, e', then we continue the search uniquely on e''. In this scenario, it may be that there is a coarsening where v lies on a 1-cell that contains e and e'.

Observe that from what we argued above, we can define an equivalence relation on the set of mutual edges. Two edges e and e' are equivalent if when when doing the operation described above from an endpoint of e, we visit e'. The fact that this relation is reflexive and transitive is trivial. The fact that it is symmetric follows from what we do when the angle is exactly π. This equivalence relation gives us a partition of the mutual edges. Observe that if an edge e is in a candidate component, then the set of edges equivalent to e are the mutual edges of that candidate component.

Now let us consider one set of equivalent mutual edges of the partition. If those mutual edges form cycles, choose such a cycle c, and orient the edges clockwisely. For any other edge of the set, orient it towards c (breaking ties arbitrarily). Let p be an endpoint of one of the mutual edge. If p is on the boundary of the convex hull, its indegree is now 0. If it is an inner point not in the regular situation, its indegree is at most 2. If it is in the regular situation, its indegree is at most 3.

Now let us consider a set E of equivalent mutual edges, such that there is no candidate component with edge set $E' \supseteq E$. As those edges in E form no cycle, there is a point p which satisfies one of these two properties: i) in any coarsener where p is a bystander it lies on a 2-face and it has indegree at most 2, ii) it is in not of the previous case, and at least one of the edges in E incident to p is not oriented towards p. If that was not the case then E would be the mutual edges of some candidate component. We orient all mutual edges in E towards p. Once again, if p is on the boundary of the convex hull, its indegree is now 0. If it is an inner point not in the regular situation, its indegree is at most 2. If it is in the regular situation, its indegree is at most 3.

Now, any remaining set of mutual edges is the set of mutual edges in some candidate component (U, E). Indeed the set of mutual edges must form a spanning tree, the points have the right indegree, no edge in E is unoriented and all edges in E with exactly one endpoint in U are oriented towards it. We choose an arbitrary endpoint p of some edge in E, and orient all mutual edges of the set towards p. We call this point the *leader* of this candidate component. For a point q different from p, if q is on the boundary of the convex hull, its indegree is now 0. If q is an inner point not in the regular situation, its indegree is at most 2. If it is in the regular situation, its indegree is at most 3.

Observe that an unoriented edge after phase 1 is still unoriented after phase 3, and thus the unoriented edges in the orientation given by the orienting algorithm are unlocked. Moreover, the orientation is admissible. We claim that for each point on the boundary of the convex hull of indegree 1, and for each inner point of indegree 3 corresponds a candidate component. We know that candidate components are prime coarseners, and that their increment is 1, therefore they are perfect coarseners. By Lemma 5.44, perfect coarseners do not fix the same pair of edges. Thus a point of degree 4 is actually part of two coarseners. Also, if a point p in the regular situation has indegree 3, then only one of the two pairs of edges incident to p consists of edges in some candidate component. So in this situation p is the leader for exactly one perfect coarsener. By construction, it is clear that the points on the boundary of the convex hull with indegree 1 are exactly those that are leaders. Thus to each leader of indegree 1 or 3 corresponds a perfect coarsener, and for each leader of indegree 4 corresponds two perfect coarseners.

5.5.2 Proof of the coarsening Lemma

Observation 5.47.

- *Let p be in $\mathcal{P}_{3,0}$, and let E_p be the set of edges incident to p. If the edges in E_p are only locked at p, then $(\{p\}, E_p)$ is a perfect coarsener.*

- *Let p be in $\mathcal{P}_{1,2}$. If the the edges e_1 and e_2 which are straight-locked at p are only locked at p, then $(\{p\}, \{e_1, e_2\})$ is a perfect coarsener.*

- *Let p be in $\mathcal{P}^b_{0,1}$, and let e be the edge straight locked at p. If e is only locked at p, then $(\{p\}, \{e\})$ is a perfect coarsener.*

- *Let p be in $\mathcal{P}_{0,4}$, let e_1, e_2 be a pair of aligned edges incident to p, and let e_3, e_4 be the other pair of edges aligned. If e_1 and e_2 (respectively e_3 and e_4) are only locked at p, then $(\{p\}, \{e_1, e_2\})$ (respectively $(\{p\}, \{e_3, e_4\})$) is a perfect coarsener.*

Lemma 5.48. *In a triangulation, no edge is locked at both endpoints.*

Proof. If an edge e is locked at both endpoints u, v, then the angle at u and v after the removal of e is at least π. Therefore, one of the face adjacent to e is not a triangle. □

Lemma 5.49 (Coarsening lemma)**.** *Let us consider $\mathcal{S}$ be a subdivision of slack D. Let us consider the refinement $\mathcal{S}_0$ of $\mathcal{S}$ obtained by skipping all bystanders in $\mathcal{S}$. We consider the orientation on $\mathcal{S}_0$ returned by the orienting algorithm. There exist at least $n - 3 - D + 2C_0 + C_1$ perfect coarsenings of $\mathcal{S}$.*

Proof. First let us assume $D = 0$, meaning that $\mathcal{S}$ is a triangulation. By Lemma 5.48, no edge is locked at both endpoints. Thus, by Observation 5.47, for each point in $\mathcal{P}_{3,0} \cup \mathcal{P}_{1,2} \cup \mathcal{P}^b_{0,1}$ there exists a perfect coarsening, and for each point in $\mathcal{P}_{0,4}$ there are two perfect coarsenings. Thus there are at least $C_{3,0} + C_{1,2} + 2C_{0,4} + C^b_{0,1}$ perfect coarsenings obtained by isolating a perfect coarsener. Let us consider the orientation on $\mathcal{S}_0$ given by the orienting algorithm. Observe that it does nothing in the second and third phases. Thus, we have simply oriented the edges towards the endpoint where they are locked, if such exists. By Observation 5.33, this orientation is admissible. Thus Lemma 5.34 implies that there are at least $N - 3 + 2C_0 + C_1 - 2C_4 - C_3 - C^b_1$ perfect coarsenings obtained by removing an unlocked edge. We observe that the points of indegree 4 are exactly those in $\mathcal{P}_{0,4}$. Moreover, the points of indegree 3 are exactly those of $\mathcal{P}_{3,0} \cup \mathcal{P}_{1,2}$. Finally, the points of indegree 1 lying on the boundary of the convex hull are exactly those in $\mathcal{P}^b_{0,1}$. By Observation 5.46, there are also $n - N$ perfect coarsenings obtained by adding a skipped point as a bystander. In total, there are at least $n - 3 + 2C_0 + C_1$ perfect coarsenings, which is the claim.

Now let us prove the statement for $1 < D < n - 3$, assuming that it holds for any $D' < D$. Let $\mathcal{S}$ be a subdivision of slack D. If there is a bystander $\mathbf{p}_0$ in $\mathcal{S}$, then the refinement $\mathcal{S}'$ obtained by skipping $\mathbf{p}_0$ is a subdivision on $n-1$ points with $\mathsf{sl}\mathcal{S}' = D-1$. Observe that the orienting algorithm outputs the same orientation on $\mathcal{S}$ and $\mathcal{S}'$ since it starts by forgetting about all bystanders anyway. Therefore there exist $(n-1)-3-(D-1)+2C_0+C_1 = n-3-D+2C_0+C_1$ perfect coarsenings of $\mathcal{S}'$. Remark that for each of these perfect coarsenings in $\mathcal{S}'$ corresponds a perfect coarsening in $\mathcal{S}$. Therefore there are $n-3-D+2C_0+C_1$ perfect coarsenings in $\mathcal{S}$.

Finally, let us assume that there is no bystander in $\mathcal{S}$. Observe that in the orientation returned by the orienting algorithm, we have one perfect coarseners for each leader of indegree 1 or 3, and two perfect coarseners for the leaders of indegree 4. As these perfect coarseners are all distinct, we obtain $2C_4+C_3+C_1^b$ perfect coarseners. Moreover, we have $N-3+2C_0+C_1-2C_4-C_3-C_1^b$ perfect coarsenings obtained by removing an unlocked edge. With the $n-N$ perfect coarseners consisting of adding a skipped point as a bystander, we obtain $n-3-D+2C_0+C_1$ perfect coarsenings. □

5.6 Link of a triangulation

In this section, we define the *link of a triangulation*. This was already defined by Wagner and Welzl for point sets in general position [82]. We extend their definition to point configurations in any dimension. We write the definition in this section instead of Section 5.3 because all results we state about it are only proven for point sets in the plane. Before giving the definition of the link, we state the following lemma about compatible almost-triangulations. Recall that in Lemma 5.25 we showed that for a skipped point p in a triangulation $\mathcal{T}$, the almost-triangulation that coarsens $\mathcal{T}$ where p is a bystander is compatible with all other almost-triangulations that coarsen $\mathcal{T}$. We show here a result with a similar flavour.

Lemma 5.50. *Let $\mathcal{T}$ be a triangulation, let p and q be two distinct involved points in $\mathcal{T}$, and let $\mathcal{S}_1$, respectively $\mathcal{S}_2$, be an almost-triangulation that coarsens $\mathcal{T}$ where p is a bystander, respectively where q is a bystander. We have that $\mathcal{S}_1$ and $\mathcal{S}_2$ are compatible.*

Proof. If there is no cell of $\mathcal{T}$ that contains both p and q, it is easy to observe that $\mathcal{S}_1$ and $\mathcal{S}_2$ are compatible. Indeed, flipping from $\mathcal{T}$ to make p a bystander has then no impact on the possibility of flipping q. If there is a cell that contains both p and q, then there exists a 1-cell B that contains both p and q, because $\mathcal{T}$ is a triangulation. Now observe that either p and q are of degree 4, or one has degree 4 and the other degree 3. In the first case, there are three 1-cells which contain p or q such that all their points are aligned. In this case we merge those three 1-cells and remove the other 1-cells that contain p or q to obtain a subdivision $\mathcal{S}_3$ where all 2-cells are of size 3, and one 1-cell contains two bystanders: p and q. In the second case, by deleting all 1-cells and that contain either p or q we obtain a subdivision $\mathcal{S}_3$ of slack 2 where all 2-cells have size 3 and one 2-cell contains two bystanders: p and q. We conclude by observing that $\mathcal{S}_3$ coarsens $\mathcal{S}_1$ and $\mathcal{S}_2$. □

5.6.1 Connectivity of the link

Definition 5.51 (Link)**.** For a triangulation $\mathcal{T}$ the *link* of $\mathcal{T}$, denoted by $\mathsf{Lk}\mathcal{T}$, is the edge-weighted graph with vertex set being the set of all almost-triangulations that coarsen $\mathcal{T}$, and there is an edge between two almost-triangulations if they are compatible. The weight of an edge $\{\mathcal{S}_1, \mathcal{S}_2\}$ is $|\mathbf{T}_{ref}\langle\mathcal{S}_{1,2}\rangle| - 2$ (which is between 1 and 3 for point sets in the plane).

In [82], Wagner and Welzl show that for a triangulation $\mathcal{T}$ on a point set with no three points on a line, $\mathsf{Lk}\mathcal{T}$ is $(n-4)$-connected. First, they argue with Lemma 5.49 that $\mathsf{Lk}\mathcal{T}$ has at least $n-3$ vertices, and that each has degree at least $n-4$ as we ourselves detail later. To show $(n-4)$-connectivity of $\mathsf{Lk}\mathcal{T}$, they show that

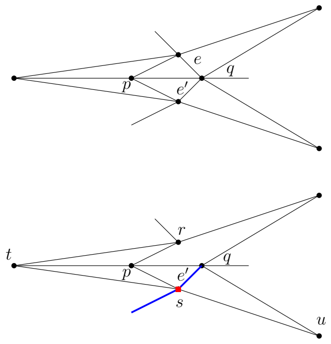

Figure 5.9: A triangulation $\mathcal{T}$ such that the complement of $\mathsf{Lk}\mathcal{T}$ contains a 4-cycle. Three of the corresponding almost-triangulations are depicted, the fourth one being symmetric to the second. In red are points of indegree at most 1 in the orientation given by the orienting algorithm. In blue are unlocked edges incident to red points.

under the assumption that no three points are on a line, there is no 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$. They conclude with the following lemma:

Lemma 5.52 ([82])**.** *Let G be a graph with its complement having no cycle of length 4, meaning that for any sequence (x_1, x_2, x_3, x_4) of four distinct vertices in G, there exists an index $i \in \{1, 2, 3, 4\}$ with $\{x_i, x_{i+1 \bmod 4}\}$ being an edge in G. Then G is δ-vertex connected, where δ is the minimum vertex degree in G.*

Interestingly, we show that when we remove the assumption that no three points are on a line, there may be 4-cycles in the complement of $\mathsf{Lk}\mathcal{T}$, as depicted in Figure 5.9. We have one triangulation $\mathcal{T}$, with some vertices p, q and some edges e, e'. The almost-triangulation where e is missing is denoted by $\mathcal{S}_1$, the one where p is a bystander is denoted by $\mathcal{S}_2$, the one where e' is missing is denoted by $\mathcal{S}_3$ and the one where $\{p, q\}$ is missing is denoted by $\mathcal{S}_4$. We list here the important properties of this triangulation. There are two aligned edges incident to p, one of them being $\{p, q\}$. Let us denote by r the endpoint of e which is not q. Removing e and the edges incident to p that are not aligned creates an angle larger than π at r. Moreover r has degree at least 5. The same holds with e' and the edges incident to p which are not aligned. This implies that we have $\mathcal{S}_1 \not\sim \mathcal{S}_2$ and $\mathcal{S}_2 \not\sim \mathcal{S}_3$. Removing e and $\{p, q\}$ creates an angle larger than π at q, and the same holds with e' and $\{p, q\}$. Thus we have $\mathcal{S}_1 \not\sim \mathcal{S}_4$ and $\mathcal{S}_3 \not\sim \mathcal{S}_4$. We have shown that $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4)$ is a sequence of almost-triangulations that form a 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$. Thus we cannot use Lemma 5.52 to show the $(n-4)$-connectivity of $\mathsf{Lk}\mathcal{T}$. We circumvent it by showing the following lemma.

Lemma 5.53. *Let $\mathcal{T}$ be a triangulation that contains the situation depicted in Figure 5.9. Let $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4$ be almost-triangulations as defined above. For any $i \in \{1, 2, 3, 4\}$, in the orientation given by the orienting algorithm on $\mathcal{S}_i$ (when forgetting about possible bystanders), there is an inner point of indegree at most* 1.

Proof. Let us consider $\mathcal{S}_2$, which is the easiest case. The blue edge is unlocked, and therefore unoriented. Observe that there must be at least two other unlocked edges incident to q, and we can even choose them such that no two are aligned. We claim that every other edge incident to q is not oriented towards q, and not mutual. Observe that the angle between any two blue edges is smaller than π. Therefore to create an angle of at least π, one would need to remove an unoriented edge. But this never happens during the second phase of the orienting algorithm, because an unoriented edge cannot be in a candidate component. Thus q has indegree 0 in the orientation.

Now let us consider $\mathcal{S}_1$. The same holds with $\mathcal{S}_3$ by symmetry. Let s denote the endpoint of e' which is not q. Observe that any coarsening of $\mathcal{S}_1$ where p is a bystander has slack at least 3, because r would also be a bystander, and then not all 2-cells can be triangle or even more points need to be bystanders. Let us denote by t the fourth neighbour of p. Observe that if t is a bystander in a coarsening of $\mathcal{S}_1$, then so is p. There are at least two unlocked edges incident to s, as shown in Figure 5.9. Let us denote by u the common neighbour of q and s depicted on the figure which is not p. Observe that after the first phase of the orienting algorithm, s has indegree at most 1, and if it is 1, then it must be from the edge $\{u, s\}$ or the edge $\{t, s\}$. Now observe that on the side of these two unlocked edges where p lies, no edge can become mutual. This is because p and t cannot be in a candidate component, as otherwise the increment would be larger than 1. Observe that u must be in any candidate component whose removal would create an angle of at least π at s. Thus only one edge incident to s can become mutual during the second phase, if $\{u, s\}$ was not already oriented towards s after the first phase. If $\{t, s\}$ is oriented towards s then no edge can become mutual.

Finally, let us consider $\mathcal{S}_4$. We claim that r and s have indegree at most 1. It is sufficient to argue for r. Observe that there exists a perfect coarsener consisting of p and the edges incident to p. Therefore the orienting algorithm will not change anything to the orientation of the edges incident to p. For this reason, t cannot be a point in any candidate component, as the edge $\{t, p\}$ is oriented towards p. Moreover q must be incident to at least one unlocked edge. Thus q cannot be in any candidate component. Let us denote by v the point adjacent to q and r in Figure 5.9. Observe that if the removal of the edges of a candidate component creates an angle of at least π at r, then v is part of this candidate component. This implies that either the edge $\{v, r\}$ is oriented towards r after the first phase, or no edge is oriented towards r after the first phase, and at most one can become mutual during the second. In both cases r has indegree at most 1 after the second phase. □

Lemma 5.54. *Let $\mathcal{T}$ be a triangulation, and let $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4)$ be*

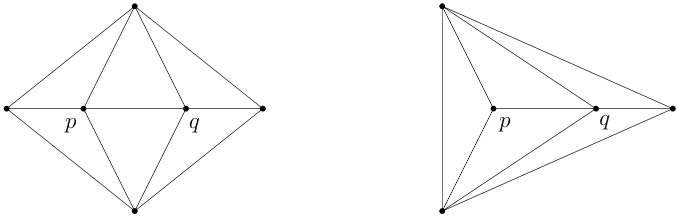

Figure 5.10: In both cases, there exist almost-triangulations where p or q is a bystander. However $\{p, q\}$ is not the intersection of two 2-cells of size 4 from two almost-triangulations as defined in the proof of Lemma 5.54.

a sequence of almost-triangulations coarsening $\mathcal{T}$, such that for any $i \in \{1, 2, 3, 4\}$, we have $\mathcal{S}_i \not\sim \mathcal{S}_{i+1 \bmod 4}$. In this situation, for any $i \in \{1, 2, 3, 4\}$, $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4)$ is the unique 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$ that contains $\mathcal{S}_i$. Moreover, in the orientation given by the orienting algorithm on $\mathcal{S}_i$ (when forgetting about possible bystanders), there is an inner point of indegree at most 1.

Proof. Let p be a point that is skipped in $\mathcal{T}$, and let $\mathcal{S}$ be the almost-triangulation identical to $\mathcal{T}$ except that p is a bystander. By Lemma 5.25, $\mathcal{S}$ is compatible with every other almost-triangulation that coarsens $\mathcal{T}$. So let us assume that no almost-triangulation in $\{\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4\}$ is identical to $\mathcal{T}$ except for a skipped point which is now a bystander. Moreover by Lemma 5.50, if $\mathcal{S}$ and $\mathcal{S}'$ are two triangulations where respectively p and q are bystanders, where p and q are involved in $\mathcal{T}$, then $\mathcal{S}$ and $\mathcal{S}'$ are compatible. Thus, we assume that no two consecutive almost-triangulations in $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4)$ are of this sort. Without loss of generality, let us assume that $\mathcal{S}_1$ and $\mathcal{S}_3$ are identical to $\mathcal{T}$ except for one 1-cell B which has disappeared, and the two 2-cells that had B has a face have now merged.

Let B_1, respectively B_3, be the 2-cell of size 4 in $\mathcal{S}_1$, respectively $\mathcal{S}_3$. Let $\mathcal{S}'$ be an almost-triangulation that coarsens $\mathcal{T}$, with $\mathcal{S}_1 \not\sim \mathcal{S}'$. If

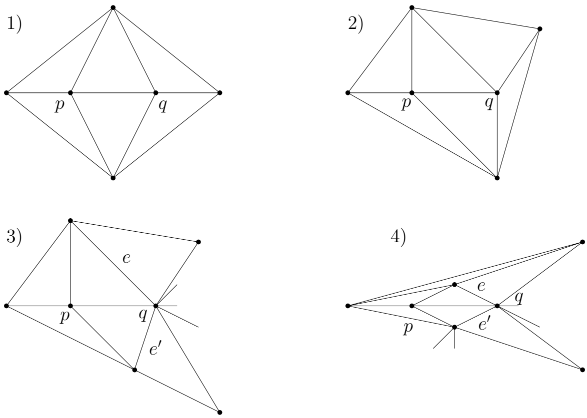

Figure 5.11: Four cases which do not satisfy the conditions of Lemma 5.54. In the two first, $\{p, q\}$ is not the intersection of two 2-cells of size 4. In the two last, there is no cycle of length 4 in the complement of $\mathsf{Lk}\mathcal{T}$.

there is a 2-cell B' of size 4 in $\mathcal{S}'$, then observe that $B_1 \cap B'$ is not empty. Moreover, the 1-cell of $\mathcal{T}$ which is not in B' must be a face of B_1. Likewise, if p is an involved point in $\mathcal{T}$ and a bystander in $\mathcal{S}'$, then p is in B_1.

We first show that $B_1 \cap B_3 = \{p, q\}$, for some points p and q. Let us assume for a contradiction that this is not the case. We show that there is at most one almost-triangulation $\mathcal{S}'$ which coarsens $\mathcal{T}$, such that $\mathcal{S}' \not\sim \mathcal{S}_1$ and $\mathcal{S}' \not\sim \mathcal{S}_3$. Therefore one of $\mathcal{S}_2, \mathcal{S}_4$ has to be compatible with $\mathcal{S}_1$ or $\mathcal{S}_3$, which is a contradiction. First, if $B_1 \cap B_3$ contains at most one element we are done. Now, as B_1 and B_3 are different, we have that $B_1 \cap B_3$ contains three points. Let us denote by e the edge in $\mathcal{T}$ that is removed in $\mathcal{S}_1$, and let e' be the edge in $\mathcal{T}$ that is removed in $\mathcal{S}_3$. Only one of the three points in $B_1 \cap B_3$ is adjacent to both e and e'. Let us denote this point by p. By construction p has degree at least 4, and is not in the regular situation. Moreover, even if some pair of edges incident to p are aligned, e and e' must be on the same side of this pair of aligned edges. So there is no almost-triangulation where p is a bystander. Let q and r denote the other points in $B_1 \cap B_3$. They have degree at least 3. If they are on the boundary of the convex hull, then there cannot be bystanders in some almost-triangulation that coarsen $\mathcal{T}$, because no pair of edges are aligned. Thus, if one of them is a bystander in some almost-triangulation, say q which is the other endpoint of e, then q has degree 4 and is incident to another edge aligned with e. In this situation, it is clear that there is no almost-triangulation where r is a bystander, as otherwise the 2-cell incident to $\{q, r\}$ which is not $B_1 \cap B_3$ would not be a triangle. Moreover, the edge $\{q, r\}$ cannot be removed in any almost-triangulation because the angle at q after its removal would be equal to π. We have shown that there is at most one almost-triangulation that is incompatible with both $\mathcal{S}_1$ and $\mathcal{S}_3$.

Now, let us assume $B_1 \cap B_3 = \{p, q\}$. Observe that $\{p, q\}$ is a face of B_1 and B_3. Indeed, 2-cells are convex, and thus if they share two elements, those form a 1-cell. We claim that there cannot be an almost-triangulation that refines $\mathcal{T}$ where p is a bystander and one where q is a bystander. Indeed, if that would be the case, then not both p and q would be of degree 3 in $\mathcal{T}$ and the situation would be one of the two cases in Figure 5.10. Observe that in both cases $\{p, q\}$ cannot be the intersection of the two convex 2-cells B_1 and B_3. Finally, we observe that there can be an almost-triangulation that refines $\mathcal{T}$ where p is a bystander and one where $\{p, q\}$ is missing. Without loss of generality, we assume that the one where p is a bystander is $\mathcal{S}_2$ and the one where $\{p, q\}$ is missing is $\mathcal{S}_4$. Observe that if there is an almost-triangulation where $\{p, q\}$ is missing, then at most one of p and q can have degree at most 4, and none of them can be of degree 3. Therefore p has degree 4. Let e denote the edge removed in B_1 and e' denote the edge removed in B_3. We show that none of the situations depicted in Figure 5.11 can happen. Cases 1) and 2) are not possible because $\{p, q\}$ cannot be the intersection of the two convex 2-cells B_1 and B_3. Case 3) is not possible because there is a coarsening of slack 2 where e is missing and p is a bystander, thus $\mathcal{S}_1$ and $\mathcal{S}_2$ would be compatible. In this coarsening, we have one bystander p, and a 2-cell of size 4, obtained by removing e and one edge incident to p. There is a second issue in Case 3), as $\mathcal{S}_3$ and $\mathcal{S}_2$ are compatible too. Indeed removing e' and the two edges incident to p which are not aligned gives a coarsening with two bystanders (p and the endpoint of e' which is not q) and all 2-cells are of size 3. Thus it has slack 2. If we are not in one of the previous cases, then removing e and the edges incident to p creates an angle larger than π at the endpoint which is not q, and the same holds with e'. Let us denote by r the other endpoint of e. Observe that this point cannot be of degree 4 as in Case 4), for in this case $\mathcal{S}_1$ and $\mathcal{S}_2$ would be compatible. Indeed, by removing the edges incident to r as well as the other edge incident to p which is not one of the aligned edges, it creates a coarsening with two bystanders (p and r) and all 2-cells are of size 3. Hence this coarsening has slack 2. A second issue with Case 4) is that removing the edges e and $\{p, q\}$ creates an angle smaller than π at q. Thus $\mathcal{S}_1$ and $\mathcal{S}_4$ are also compatible. Finally, we observe that we must be in the situation depicted in Figure 5.9. We conclude with Lemma 5.53 that we have for all subdivisions an inner point with indegree at most 1. It can be easily observed that for any $i \in \{1, 2, 3, 4\}$, there is only one 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$ that contains $\mathcal{S}_i$. □

Lemma 5.55. *Given three triangulations $\mathcal{T}_1$, $\mathcal{T}_2$ and $\mathcal{T}_3$ and a pair of almost-triangulations $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $\mathcal{S}_1$ coarsens $\mathcal{T}_1$ and $\mathcal{T}_3$ and such that $\mathcal{S}_2$ coarsens $\mathcal{T}_3$ and $\mathcal{T}_2$, every $\mathcal{S}_1$-$\mathcal{S}_2$-path of weight w in $\mathsf{Lk}\mathcal{T}_3$ induces a $\mathcal{T}_3$-avoiding $\mathcal{T}_1$-$\mathcal{T}_2$-path of length w in the bistellar flip graph. Internally vertex disjoint $\mathcal{S}_1$-$\mathcal{S}_2$-paths in $\mathsf{Lk}\mathcal{T}_3$ induce internally vertex disjoint $\mathcal{T}_1$-$\mathcal{T}_2$-paths in the bistellar flip graph.*

Proof. Let us consider a $\mathcal{S}_1$-$\mathcal{S}_2$-path in $\mathsf{Lk}\mathcal{T}_3$. Let $\{\mathcal{S}, \mathcal{S}'\}$ be an edge of this path. By assumption, there exists a subdivision of slack 2 $\mathcal{S}''$

that coarsen $\mathcal{S}$ and $\mathcal{S}'$. By Lemma 5.23, there exists a path in the bistellar flip graph from $\mathcal{T}_3[\mathcal{S}]$ to $\mathcal{T}_3[\mathcal{S}']$ of length between 1 and 3, which does not go through $\mathcal{T}_3$. The internal vertices of this path are the triangulations in $\mathbf{T}_{ref}\langle\mathcal{S}''\rangle \setminus \{\mathcal{T}_3[\mathcal{S}], \mathcal{T}_3, \mathcal{T}_3[\mathcal{S}']\}$. Those vertices are at distance 2 from $\mathcal{T}_3$ in the bistellar flip graph, and $\mathcal{T}_3[\mathcal{S}]$ and $\mathcal{T}_3[\mathcal{S}']$ are both at distance 1. We concatenate all these paths to obtain one from $\mathcal{T}_3[\mathcal{S}_1] = \mathcal{T}_1$ to $\mathcal{T}_3[\mathcal{S}_2] = \mathcal{T}_2$.

Let us consider a set of internally vertex disjoints $\mathcal{S}_1$-$\mathcal{S}_2$-path in $\mathsf{Lk}\mathcal{T}_3$. Let $\mathcal{S}'$ be an internal vertex of one of these paths. In the corresponding path of the bistellar flip graph, $\mathcal{S}'$ corresponds to $\mathcal{T}_3[\mathcal{S}']$. This implies that internal vertices at distance 1 from $\mathcal{T}_3$ in the paths on the bistellar flip graphs are all distinct. It remains to show that vertices at distance 2 are also all distinct. Suppose we have $\mathcal{T}$ and $\mathcal{T}'$, both at distance 2 from $\mathcal{T}_3$, that appear in some paths from $\mathcal{T}_1$ to $\mathcal{T}_2$. Thus there exists $\mathcal{S}$ and $\mathcal{S}'$ of slack 2 such that $\mathcal{S}$ coarsen $\mathcal{T}$ and $\mathcal{T}_3$, and $\mathcal{S}'$ coarsen $\mathcal{T}'$ and $\mathcal{T}_3$. By Lemma 5.21, $\mathbf{T}_{ref}\langle\mathcal{S}\rangle \cap \mathbf{T}_{ref}\langle\mathcal{S}'\rangle$ is either equal to $\mathbf{T}_{ref}\langle\mathcal{S}_0\rangle$ for some almost-triangulation $\mathcal{S}_0$ or we have $\mathcal{S} = \mathcal{S}'$. In the first case, $\mathbf{T}_{ref}\langle\mathcal{S}\rangle$ and $\mathbf{T}_{ref}\langle\mathcal{S}'\rangle$ cannot share a triangulation at distance 2 from $\mathcal{T}_3$, thus $\mathcal{T}$ and $\mathcal{T}'$ cannot be the same. If $\mathcal{S} = \mathcal{S}'$, then they correspond to the same path in $\mathsf{Lk}\mathcal{T}_3$. □

Lemma 5.56. *For a triangulation $\mathcal{T}$, the link $\mathsf{Lk}\mathcal{T}$ is $(n-4)$-vertex connected.*

Proof. First, $\mathsf{Lk}\mathcal{T}$ has at least $n-3$ vertices because by Lemma 5.49, there are at least $n-3$ almost-triangulation that coarsen $\mathcal{T}$. Moreover, we have observed that no two of these almost-triangulations leads to the same neighbour of $\mathcal{T}$ in the bistellar flip graph.

If $\mathcal{S}_1$ is a vertex in $\mathsf{Lk}\mathcal{T}$, then by Lemma 5.17 $\mathcal{S}_1$ has slack 1. Thus $\mathcal{S}_1$ has $(n-4+2C_0+C_1)$ perfect coarseners by Lemma 5.49, where C_i denotes the number of inner point with indegree i in the orientation given by the orienting algorithm on $\mathcal{S}_1$ (when forgetting about possible bystanders). Each subdivision obtained by isolating such a perfect coarsener is equal to $\mathcal{S}_{1,2}$ for some almost-triangulation $\mathcal{S}_2$ with $\mathcal{S}_1 \diamond \mathcal{S}_2$, as implied by Lemma 5.24, and this $\mathcal{S}_{1,2}$ is uniquely defined by the perfect coarsener. This implies that $\mathcal{S}_1$ has minimum degree at least $(n-4+2C_0+C_1)$ in $\mathsf{Lk}\mathcal{T}$.

Now let us consider two almost-triangulations $\mathcal{S}_1$ and $\mathcal{S}_2$ that coarsen $\mathcal{T}$, such that they are at distance 2 in $\mathsf{Lk}\mathcal{T}$. By Lemma 5.14, to show that $\mathsf{Lk}\mathcal{T}$ is $(n-4)$-connected, it is sufficient to show that there exist $(n-4)$ internally vertex disjoint $\mathcal{S}_1$-$\mathcal{S}_2$-paths. First let us assume that no 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$ contains $\mathcal{S}_1$ and $\mathcal{S}_2$. Both almost-triangulations have at least $(n-4)$ neighbours in $\mathsf{Lk}\mathcal{T}$. Let us denote by ℓ the number of common neighbours they have. If $\ell \geq n-4$ we are done. Otherwise there are $(n-4-\ell)$ almost-triangulations compatible with $\mathcal{S}_1$ but not with $\mathcal{S}_2$, and $(n-4-\ell)$ almost-triangulations compatible with $\mathcal{S}_2$ but not with $\mathcal{S}_1$. Let us consider one almost-triangulation $\mathcal{S}_3$ and one almost-triangulation $\mathcal{S}_4$ from each group. We have $\mathcal{S}_1 \not\diamond \mathcal{S}_2$, $\mathcal{S}_3 \not\diamond \mathcal{S}_2$ and $\mathcal{S}_4 \not\diamond \mathcal{S}_1$. As there is no 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$ that contains $\mathcal{S}_1$ and $\mathcal{S}_2$, we infer $\mathcal{S}_3 \diamond \mathcal{S}_4$. Thus, we have found ℓ paths of length 2 and $(n-4-\ell)$-paths of length 3 that are pairwise internally vertex disjoint.

Finally let us assume that there exists a 4-cycle in the complement of $\mathsf{Lk}\mathcal{T}$ that contains $\mathcal{S}_1$ and $\mathcal{S}_2$. By Lemma 5.54, there exist $\mathcal{S}_3$ and $\mathcal{S}_4$ such that $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4)$ forms a cycle in the complement of $\mathsf{Lk}\mathcal{T}$, and this 4-cycle is the unique one that contains $\mathcal{S}_1$ and $\mathcal{S}_2$. Moreover, Lemma 5.54 implies that the orienting algorithm on $\mathcal{S}_1$ or $\mathcal{S}_2$ returns an orientation with at least one inner point with indegree at most 1. This implies that $\mathcal{S}_1$ and $\mathcal{S}_2$ have degree at least $(n-3)$ in $\mathsf{Lk}\mathcal{T}$. Let ℓ denote the number of common neighbours of $\mathcal{S}_1$ and $\mathcal{S}_2$. If $\ell \geq (n-4)$ we are done. Otherwise we can argue as above. It might be that one of the remaining neighbours of $\mathcal{S}_1$ is not adjacent to one of the remaining neighbours of $\mathcal{S}_2$, but then those two neighbours must be $\mathcal{S}_3$ and $\mathcal{S}_4$. Thus, as there are at least $(n-3-\ell)$ remaining neighbours, we can find in total at least $(n-4)$ internally vertex disjoint $\mathcal{S}_1$-$\mathcal{S}_2$-paths. □

5.6.2 Proof of connectivity of the bistellar flip graph

We prove Theorem 5.12. This last part is exactly as in [82], but we state it as our notation is quite different from theirs, following the fact the we allow for several points to be on a line.

Proof of Theorem 5.12. If n is at most 4, we can check easily $(n-3)$-connectivity according to the definition of k-vertex connectivity. For $n \geq 5$, we can use Lemma 5.14. There are at least $(n-2)$ triangulations since the graph is non-empty and each triangulation has degree at least $n-3$ by Lemma 5.49. By Lemma 5.14, it is sufficient to show that for any pair of triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ at distance 2 in the bistellar flip graph, there are $(n-3)$ pairwise internally vertex disjoint $\mathcal{T}_1$-$\mathcal{T}_2$-paths. By assumption, there exists $\mathcal{T}_3$ such that there is a flip between $\mathcal{T}_1$ and $\mathcal{T}_3$, as well as a flip between $\mathcal{T}_3$ and $\mathcal{T}_2$. Let us denote by $\mathcal{S}_1$ the almost-triangulation corresponding to the flip between $\mathcal{T}_1$ and $\mathcal{T}_3$, and let $\mathcal{S}_2$ be the almost-triangulation corresponding to the flip between $\mathcal{T}_3$ and $\mathcal{T}_2$. By Lemma 5.56, there exists $(n-4)$ internally vertex disjoint paths between $\mathcal{T}_3[\mathcal{S}_1]$ and $\mathcal{T}_3[\mathcal{S}_2]$ in $\mathsf{Lk}\mathcal{T}_3$. Therefore Lemma 5.55 implies the existence of $(n-4)$ internally vertex disjoint paths between $\mathcal{T}_1$ and $\mathcal{T}_2$ in the bistellar flip graph, which avoid $\mathcal{T}_3$. In total, we have found $(n-3)$ internally vertex disjoint paths between $\mathcal{T}_1$ and $\mathcal{T}_2$ in the bistellar flip graph. □

Bibliography

[1] A. Aamand, M. Abrahamsen, J. B. T. Knudsen, and P. M. R. Rasmussen. Classifying convex bodies by their contact and intersection graphs. *CoRR*, abs/1902.01732, 2019. URL: `http://arxiv.org/abs/1902.01732`, `arXiv:1902.01732`.

[2] A. Agrawal, G. Guspiel, J. Madathil, S. Saurabh, and M. Zehavi. Connecting the dots (with minimum crossings). In *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA.*, pages 7:1–7:17, 2019. `doi:10.4230/LIPIcs.SoCG.2019.7`.

[3] P. Alimonti and V. Kann. Some APX-completeness results for cubic graphs. *Theor. Comput. Sci.*, 237(1-2):123–134, 2000. `doi:10.1016/S0304-3975(98)00158-3`.

[4] B. Aronov, A. Donakonda, E. Ezra, and R. Pinchasi. On pseudo-disk hypergraphs. *arXiv preprint arXiv:1802.08799*, 2018.

[5] S. Bandyapadhyay, S. Bhowmick, and K. Varadarajan. Approximation schemes for partitioning: Convex decomposition and surface approximation. In *Proceedings of the Twenty-Sixth*

Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1457–1470. SIAM, 2014. doi:10.1137/1.9781611973730.96.

[6] A. S. Barboza, C. C. de Souza, and P. J. de Rezende. Minimum convex partition of point sets. In *Proceedings of International Conference on Algorithms and Complexity*, pages 25–37. Springer, 2019. doi:/10.1007/978-3-030-17402-6_3.

[7] H. Bergold, D. Bertschinger, N. Grelier, W. Mulzer, and P. Schnider. Well-separation and hyperplane transversals in high dimensions.

[8] M. Bonamy, É. Bonnet, N. Bousquet, P. Charbit, and S. Thomassé. EPTAS for max clique on disks and unit balls. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 568–579, 2018. doi:10.1109/FOCS.2018.00060.

[9] É. Bonnet, P. Giannopoulos, E. J. Kim, P. Rzążewski, and F. Sikora. QPTAS and subexponential algorithm for maximum clique on disk graphs. In *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, pages 12:1–12:15, 2018. doi:10.4230/LIPIcs.SoCG.2018.12.

[10] É. Bonnet, N. Grelier, and T. Miltzow. Maximum clique in disk-like intersection graphs. *arXiv preprint arXiv:2003.02583*, 2020.

[11] É. Bonnet and T. Miltzow. Parameterized hardness of art gallery problems. In *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, pages 19:1–19:17, 2016. doi:10.4230/LIPIcs.ESA.2016.19.

[12] P. Bose, A. Brodnik, S. Carlsson, E. D. Demaine, R. Fleischer, A. López-Ortiz, P. Morin, and J. Ian Munro. Online routing in convex subdivisions. *International Journal of Computational*

Geometry & Applications, 12(04):283–295, 2002. doi:10.1142/S021819590200089X.

[13] H. Breu and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Comput. Geom.*, 9(1-2):3–24, 1998. doi:10.1016/S0925-7721(97)00014-X.

[14] V. E. Brimkov, K. Junosza-Szaniawski, S. Kafer, J. Kratochvíl, M. Pergel, P. Rzążewski, M. Szczepankiewicz, and J. Terhaar. Homothetic polygons and beyond: Maximal cliques in intersection graphs. *Discrete Applied Mathematics*, 247:263–277, 2018. doi:10.1016/j.dam.2018.03.046.

[15] A. Butman, D. Hermelin, M. Lewenstein, and D. Rawitz. Optimization problems in multiple-interval graphs. *ACM Trans. Algorithms*, 6(2):40:1–40:18, 2010. doi:10.1145/1721837.1721856.

[16] S. Cabello, J. Cardinal, and S. Langerman. The clique problem in ray intersection graphs. *Discrete & Computational Geometry*, 50(3):771–783, 2013. doi:10.1007/s00454-013-9538-5.

[17] H. Cambazard and N. Catusse. An integer programming formulation using convex polygons for the convex partition problem. In *37th International Symposium on Computational Geometry (SoCG 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.SoCG.2021.20.

[18] P. L. Chew. Constrained delaunay triangulations. *Algorithmica*, 4(1-4):97–108, 1989. doi:10.1007/BF01553881.

[19] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1-3):165–177, 1990. doi:10.1016/0012-365X(90)90358-O.

[20] J. Cleve, N. Grelier, K. Knorr, M. Löffler, W. Mulzer, and D. Perz. Nearest-neighbor decompositions of drawings. 37th Eu-

ropean Workshop on Computational Geometry (EuroCG 2021), 2021.

[21] F. d'Amore and P. G. Franciosa. On the optimal binary plane partition for sets of isothetic rectangles. *Information Processing Letters*, 44(5):255–259, 1992. doi:10.1016/0020-0190(92)90210-M.

[22] J. De Loera, J. Rambau, and F. Santos. *Triangulations: structures for algorithms and applications*, volume 25. Springer Science & Business Media, 2010.

[23] J. A. De Loera, F. Santos, and J. Urrutia. The number of geometric bistellar neighbors of a triangulation. *Discrete & Computational Geometry*, 21(1):131–142, 1999.

[24] E. Demaine, S. Fekete, P. Keldenich, D. Krupke, and J. S. B. Mitchell. CG:SHOP 2020. https://cgshop.ibr.cs.tu-bs.de/competition/cg-shop-2020. Accessed: 12/02/2020.

[25] A. Dumitrescu, S. Har-Peled, and C. D. Tóth. Minimum convex partitions and maximum empty polytopes. In *Proceedings of Scandinavian Workshop on Algorithm Theory*, pages 213–224. Springer, 2012. doi:10.1007/978-3-642-31155-0_19.

[26] A. Dumitrescu, J. S. B. Mitchell, and M. Sharir. Binary space partitions for axis-parallel segments, rectangles, and hyperrectangles. *Discrete & Computational Geometry*, 31(2):207–227, 2004. doi:10.1007/s00454-003-0729-3.

[27] Z. Dvořák and J. Pekárek. Induced odd cycle packing number, independent sets, and chromatic number. *arXiv preprint arXiv:2001.02411*, 2020.

[28] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.

[29] T. Fevens, H. Meijer, and D. Rappaport. Minimum convex partition of a constrained point set. *Discrete Applied Mathematics*, 109(1-2):95–107, 2001. doi:10.1016/S0166-218X(00)00237-7.

[30] M. C. Francis, D. Gonçalves, and P. Ochem. The Maximum Clique Problem in Multiple Interval Graphs. *Algorithmica*, 71(4):812–836, 2015. doi:10.1007/s00453-013-9828-6.

[31] J. García-López and C. M. Nicolás. Planar point sets with large minimum convex decompositions. *Graphs and Combinatorics*, 29(5):1347–1353, 2013. doi:10.1007/s00373-012-1181-z.

[32] M. R. Garey and D. S. Johnson. The rectilinear steiner tree problem is np-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977.

[33] M. Grantson and C. Levcopoulos. A fixed parameter algorithm for the minimum number convex partition problem. In *Japanese Conference on Discrete and Computational Geometry*, pages 83–94. Springer, 2004. doi:10.1007/11589440_9.

[34] N. Grelier. Computing a maximum clique in geometric superclasses of disk graphs. In *International Computing and Combinatorics Conference*, pages 299–310. Springer, 2020.

[35] N. Grelier. Hardness and Approximation of Minimum Convex Partition. 224:45:1–45:15, 2022. URL: https://drops.dagstuhl.de/opus/volltexte/2022/16053, doi:10.4230/LIPIcs.SoCG.2022.45.

[36] N. Grelier, S. Ilchi, T. Miltzow, and S. Smorodinsky. On the vc-dimension of half-spaces with respect to convex sets. *Discrete Mathematics & Theoretical Computer Science*, 2021.

[37] N. Grelier, R. d. Joannis de Verclos, R. J. Kang, and F. Pirot. Approximate strong edge-colouring of unit disk graphs. In *Inter-*

national Workshop on Approximation and Online Algorithms, pages 154–169. Springer, 2019.

[38] L. J. Guibas, M. H. Overmars, and J.-M. Robert. The exact fitting problem in higher dimensions. *Computational geometry*, 6(4):215–230, 1996. doi:10.1016/0925-7721(95)00020-8.

[39] U. I. Gupta, D.-T. Lee, and J.-T. Leung. Efficient algorithms for interval graphs and circular-arc graphs. *Networks*, 12(4):459–467, 1982.

[40] G. Guspiel. Complexity of finding perfect bipartite matchings minimizing the number of intersecting edges. *CoRR*, abs/1709.06805, 2017. URL: http://arxiv.org/abs/1709.06805, arXiv:1709.06805.

[41] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098.

[42] T. Hertli. 3-SAT faster and simpler - unique-SAT bounds for PPSZ hold in general. *SIAM J. Comput.*, 43(2):718–729, 2014. doi:10.1137/120868177.

[43] H. Imai and T. Asano. Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. *Journal of algorithms*, 4(4):310–323, 1983.

[44] R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.

[45] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, Dec. 2001.

[46] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter. Hamilton paths in grid graphs. *SIAM Journal on Computing*, 11(4):676–686, 1982.

[47] D. S. Johnson. Approximation algorithms for combinatorial problems. In *Proceedings of the fifth annual ACM Symposium on Theory of Computing*, pages 38–49, 1973. doi:10.1145/800125.804034.

[48] D. S. Johnson. The NP-completeness column: an ongoing guide. *Journal of algorithms*, 6(3):434–451, 1985.

[49] R. J. Kang and T. Müller. Sphere and dot product representations of graphs. *Discrete & Computational Geometry*, 47(3):548–568, 2012.

[50] M. J. Keil. Decomposing a polygon into simpler components. *SIAM Journal on Computing*, 14(4):799–817, 1985. doi:10.1137/0214056.

[51] S.-J. Kim, A. Kostochka, and K. Nakprasit. On the chromatic number of intersection graphs of convex sets in the plane. *the electronic journal of combinatorics*, 11(1):52, 2004.

[52] C. Knauer and A. Spillner. Approximation algorithms for the minimum convex partition problem. In *Proceedings of Scandinavian Workshop on Algorithm Theory*, pages 232–241. Springer, 2006. doi:10.1007/11785293_23.

[53] J. Kratochvíl and J. Nešetřil. Independent set and clique problems in intersection-defined classes of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 31(1):85–93, 1990. Available at http://eudml.org/doc/17810.

[54] A. V. Kumar, S. Arya, and H. Ramesh. Hardness of set cover with intersection 1. In *Proceedings of International Colloquium on Automata, Languages, and Programming*, pages 624–635. Springer, 2000. doi:10.1007/3-540-45022-X_53.

[55] S. Langerman and P. Morin. Covering things with things. *Discrete & Computational Geometry*, 33(4):717–729, 2005. doi:10.1007/s00454-004-1108-4.

[56] C. L. Lawson. Transforming triangulations. *Discrete mathematics*, 3(4):365–372, 1972.

[57] D. Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982. doi:10.1137/0211025.

[58] A. Lingas. The power of non-rectilinear holes. In *Proceedings of International Colloquium on Automata, Languages, and Programming*, pages 369–383. Springer, 1982. doi:10.1007/BFb0012784.

[59] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975. doi:10.1016/0012-365X(75)90058-8.

[60] D. Marx. Parameterized complexity of independence and domination on geometric graphs. In *International Workshop on Parameterized and Exact Computation*, pages 154–165. Springer, 2006. doi:10.1007/11847250_14.

[61] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 865–877, 2015. doi:10.1007/978-3-662-48350-3_72.

[62] M. Middendorf and F. Pfeiffer. The max clique problem in classes of string-graphs. *Discrete Mathematics*, 108(1-3):365–372, 1992. doi:10.1016/0012-365X(92)90688-C.

[63] J. S. B. Mitchell. Approximation algorithms for geometric separation problems. *Technical report, Dept. of Applied Math. and Statistics, State U. of New York at Stony Brook*, 1993.

Available at https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.50.7089&rep=rep1&type=pdf.

[64] D. Moshkovitz and R. Raz. Sub-constant error probabilistically checkable proof of almost-linear size. *Computational Complexity*, 19(3):367–422, 2010. doi:10.1007/s00037-009-0278-0.

[65] V. Neumann-Lara, E. Rivera-Campo, and J. Urrutia. A note on convex decompositions of a set of points in the plane. *Graphs and Combinatorics*, 20(2):223–231, 2004. doi:10.1007/s00373-004-0555-2.

[66] M. S. Paterson and F. F. Yao. Optimal binary space partitions for orthogonal objects. *Journal of algorithms*, 13(1):99–113, 1992. doi:10.1016/0196-6774(92)90007-Y.

[67] R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for k-SAT. *J. ACM*, 52(3):337–364, 2005. doi:10.1145/1066100.1066101.

[68] S. Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.

[69] F. P. Preparata and M. I. Shamos. *Computational Geometry*. Springer-Verlag, New York, 1985. doi:10.1007/978-1-4612-1098-6.

[70] V. Raghavan and J. P. Spinrad. Robust algorithms for restricted domains. *J. Algorithms*, 48(1):160–172, 2003. doi:10.1016/S0196-6774(03)00048-8.

[71] T. Sakai and J. Urrutia. Convex decompositions of point sets in the plane. *arXiv preprint arXiv:1909.06105*, 2019.

[72] F. Santos. A point set whose space of triangulations is disconnected. *Journal of the American Mathematical Society*, 13(3):611–637, 2000.

[73] F. Santos. Non-connected toric hilbert schemes. *Mathematische Annalen*, 332(3):645–665, 2005.

[74] A. Sapucaia, P. J. de Rezende, and C. C. de Souza. Solving the minimum convex partition of point sets with integer programming. *Computational Geometry*, page 101794, 2021. doi:10.1016/j.comgeo.2021.101794.

[75] A. Spillner. A fixed parameter algorithm for optimal convex partitions. *Journal of Discrete Algorithms*, 6(4):561–569, 2008. doi:10.1016/j.jda.2008.07.002.

[76] S. Tippenhauer. On planar 3-SAT and its variants. *Master Thesis, Fachbereich Mathematik und Informatik der Freien Universitat Berlin*, 2016.

[77] C. D. Tóth. A note on binary plane partitions. In *Proceedings of the seventeenth annual symposium on Computational geometry*, pages 151–156, 2001. doi:10.1145/378583.378652.

[78] C. D. Tóth. Binary space partitions for line segments with a limited number of directions. *SIAM Journal on Computing*, 32(2):307–325, 2003. doi:10.1137/S0097539702403785.

[79] C. D. Tóth. Binary plane partitions for disjoint line segments. *Discrete & Computational Geometry*, 45(4):617–646, 2011. doi:10.1007/s00454-011-9341-0.

[80] L. G. Valiant. Universality considerations in VLSI circuits. *IEEE Transactions on Computers*, 30(2):135–140, 1981. doi:10.1109/TC.1981.6312176.

[81] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, 16:264–280, 1971.

[82] U. Wagner and E. Welzl. Connectivity of triangulation flip graphs in the plane (part II: Bistellar flips). In *36th Interna-*

tional Symposium on Computational Geometry (SoCG 2020). Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

Curriculum Vitae

Nicolas Grelier

born on the 16th of March 1995
French citizen

2018–2022	Ph.D. in Theoretical Computer Science ETH Zürich, Zürich, Switzerland
2015–2018	M.Sc. in Engineering "Diplôme d'Ingénieur" IMT Atlantique, Brest, France
2013–2015	Higher School Preparatory Classes Lycée Clemenceau, Nantes, France
2010–2013	French baccalaureate Lycée Clemenceau, Nantes, France