Diss. ETH No. 18607

Crossing-Free Configurations on Planar Point Sets

A dissertation submitted to the
SWISS FEDERAL INSTITUTE OF TECHNOLOGY
ZURICH

for the degree of
DOCTOR OF SCIENCES

presented by
ANDREAS RAZEN
Dipl. Math. ETH
born May 7, 1982
citizen of Austria

accepted on the recommendation of
Prof. Dr. Emo Welzl, examiner
Prof. Dr. Jack Snoeyink, co-examiner
Dr. Uli Wagner, co-examiner

2009

to my parents

Acknowledgments

I am deeply grateful to my supervisor Emo Welzl for his motivating and encouraging support during my time as Ph.D. student. He introduced me to the beautiful field of geometric graph theory, and with his inspiring advice he made it an unforgettable experience to learn from him. It was an honor being part of his research group Gremo, which provided both a stimulating research environment and a wonderful academic home where I met people who I will greatly miss.

My sincerest thanks go to Jack Snoeyink for working with me and giving me the opportunity to learn about applications of geometric computations in molecular biology, for reviewing my thesis and providing me with many valuable suggestions and comments which greatly helped to improve the presentation of this dissertation.

I am also indebted to Uli Wagner for working together with me on the transformation graphs of compatible spanning trees, for broadening my horizon by teaching vivid courses on Fourier-analytic, topological, and algebraic methods in combinatorics, and for co-refereeing my thesis and commenting on its write-up.

I would like to thank all the current and previous members of Gremo for making my time as Ph.D. student such a pleasant experience: Robert Berke, Yves Brise, Péter Csorba, Tobias Christ, Kaspar Fischer, Bernd Gärtner, Heidi Gebauer, Joachim Giesen, Franziska Hefti, Michael Hoffmann, Martin Jaggi, Shankar Ram Lakshminarayanan, Dieter Mitsche, Robin Moser, Leo Rüst, Andrea Salow, Dominik Scheder, Eva Schuberth, Miloš Stojaković, Marek Sulovský, Tibor Szabó, Patrick Traxler, Floris Tschurr, Frans Wessendorp, and Philipp Zumstein.

In particular I am thankful to Andrea, Franziska and Floris for resolving all kinds of administrative issues; Bernd and Michael for teaching fascinating courses on computational geometry; Philipp for studying and working with me for one third of my lifetime, and for pointing me out to Emo's Satisfiability lecture and Tibor's Graph theory course back in 2003; Dominik for carefully reading drafts of my papers and this thesis, and for giving many talks in the Mittagsseminar; Tobias for almost sharing the same birthday with me and for celebrating two of these occasions on legen' – wait for it – 'dary visits to Jerusalem and Berlin; Yves and Marek for a jet-set trip to Prague and for sharing the interest in a certain TV sitcom; Martin for reading parts of this thesis and for coding in the Algorithms Lab with me; Heidi for suggestions regarding the write-up of my thesis; Patrick for his Austrian sense of humor; and Robin for repeatedly asking me how my job applications were coming along.

Last but not least my special thanks go to my family and my friends outside ETH for their help and support, without whom this project would not have been possible. Particularly, to my parents Gisela and Reinhard for always being there, to my brothers Alexander and Michael for continuously challenging me, and to my girlfriend Janine for her love and understanding.

Zurich, September 2009

Contents

Abstract

This thesis treats crossing-free geometric graphs, which are graphs defined on a given finite set of points in the plane. Their edges are understood as straight-line segments connecting corresponding endpoints without passing through any other point. For a crossing-free graph no two edges are allowed to intersect other than in a common endpoint. The following are the most important kinds of crossing-free configurations that we will encounter in our discussion: Triangulations, crossing-free partitions, spanning trees, perfect matchings, and the set of all plane graphs.

First, we consider a class of plane graphs that emerge from crossing-free partitions of the underlying point set. To be more precise, a partition of a given set of points is crossing-free if the convex hulls of the individual parts are mutually disjoint. Each partition naturally translates to a plane graph whose vertices are the given points and whose edges are the boundaries of the convex hulls of the partition classes. We ask whether convex position of the underlying point set minimizes the number of crossing-free partitions over all placements of equally many points. We answer a corresponding question in the affirmative for the number of crossing-free partitions of n points into a fixed number k of parts, where $k \in \{1, 2, 3, n-3, n-2, n-1, n\}$. In addition, we show that on at least five points the number of crossing-free partitions is not maximized in convex position. It is known that in convex position the number of crossing-free partitions into k classes equals the number of partitions into $n - k + 1$ parts. This does not hold in general, and we mention a construction for point sets with significantly more partitions into few classes than into many. Another problem we consider on point sets in convex position is the decomposition of the complete graph using geometric graphs corresponding to crossing-free partitions. We show almost tight bounds for the number of elements in a smallest possible decomposition.

Second, we treat transformation graphs of crossing-free configurations on a set of points. These are abstract graphs whose vertices are the crossing-free configurations of interest and whose edges are defined by a prescribed transformation rule. For instance, we call two configurations compatible if the union of their edge sets is again crossing-free. In the corresponding transformation graph every pair of compatible configurations is joined by an edge. As the transformation rule encodes some notion of similarity of the configurations, the diameter of the transformation graph is a natural parameter to study. For two classes of transformation graphs – compatible crossing-free spanning trees and compatible crossing-free perfect matchings – we provide constructions for placements of n points such that the diameter is $\Omega(\log n/\log\log n)$. This nearly matches the known upper bound of $O(\log n)$ in both cases. For the transformation graph of compatible spanning trees our construction yields a tight result in terms of the number k of convex layers of the given point set, i.e., the diameter is $\Omega(\log k)$ which is best-possible.

In the last part of this thesis we study algorithmic aspects of counting the number of crossing-free geometric graphs on a given set of points. We show that the total number of such crossing-free graphs can be computed with exponential speed-up compared to enumerating them. It is worth emphasizing that no similar statements are known for other prominent graph classes such as triangulations, spanning trees or perfect matchings. Another result we obtain is a lower bound on the total number of plane graphs, in the sense that there are at least $\sqrt{8}^{n-1}$ times more crossing-free graphs than triangulations on any set of n points in general position. Upper bounds on the number of plane graphs are usually derived in a straight-forward way by providing estimates on the number of triangulations and then counting all subsets of edges in each triangulation. While there exist degenerate point sets for which the estimate obtained by counting subgraphs is best-possible, we are able to improve over this trivial upper bound assuming general position of the underlying point set. In particular, we obtain the currently best upper bound of $O(343.106^n)$ on the total number of crossing-free graphs a set of n points can have.

Zusammenfassung

Gegenstand der vorliegenden Dissertation sind kreuzungsfreie geometrische Graphen. Geometrische Graphen sind Graphen, deren Knoten durch eine endliche Menge von Punkten in der Ebene definiert sind. Ihre Kanten sind durch Liniensegmente gegeben, welche entsprechende Endpunkte miteinander verbinden, ohne durch einen weiteren Punkt zu verlaufen. In einem kreuzungsfreien Graphen dürfen sich keine zwei Kanten schneiden ausser in einem ihrer Endpunkte. Die wichtigsten kreuzungsfreien Konfigurationen, welche in dieser Arbeit behandelt werden, sind die folgenden: Triangulierungen, kreuzungsfreie Partitionen, Spannbäume, perfekte Paarungen und die Menge aller planarer Graphen.

Zunächst betrachten wir eine Klasse von planaren Graphen, die aus kreuzungsfreien Partitionen der zugrundeliegenden Punktmenge hervorgehen. Eine Partition einer gegebenen Menge von Punkten heisst kreuzungsfrei, falls die konvexen Hüllen der einzelnen Klassen paarweise disjunkt sind. Solche Partitionen lassen sich auf natürliche Weise als planare Graphen auf der Punktmenge interpretieren, deren Kanten aus den Rändern der konvexen Hüllen der einzelnen Partitionsklassen bestehen. Wir untersuchen die Frage, ob die konvexe Lage unter allen Anordnungen gleich vieler Punkte die Anzahl der kreuzungsfreien Partitionen minimiert. Für die Zahl der kreuzungsfreien Partition von n Punkten in k Klassen, wobei $k \in \{1, 2, 3, n-3, n-2, n-1, n\}$, können wir die entsprechende Frage positiv beantworten. Sind mindestens fünf Punkte gegeben, zeigen wir ausserdem, dass die konvexe Lage die Anzahl der kreuzungsfreien Partitionen nicht maximiert. Es ist bekannt, dass in konvexer Lage die Zahl der kreuzungsfreien Partitionen in k und in $n - k + 1$ Klassen übereinstimmen. Dies gilt im Allgemeinen nicht, und wir konstruieren Punktmengen, die für kleines k signifikant mehr Partitionen in k als in $n - k + 1$ Klassen erlauben. Ein weiteres Problem, welches wir auf Punktmengen in konvexer Lage betrachten, ist das Zerlegen des vollständigen Graphen in planare Graphen, die aus kreuzungsfreien Partitionen entstehen. Für die Anzahl der Elemente einer kleinstmöglichen Zerlegung beweisen wir fast scharfe Schranken.

Ebenfalls behandeln wir Transformationsgraphen kreuzungsfreier Konfigurationen, die auf Punktmengen definiert sind. Transformationsgraphen sind abstrakte Graphen, deren Knoten die kreuzungsfreien Konfigurationen darstellen, und deren Kanten durch eine vorgegebene Transformationsregel bestimmt sind. Eine solche Regel ist zum Beispiel durch Kompatibilität gegeben, wobei zwei Konfigurationen kompatibel heissen, falls die Vereinigung ihrer Kantenmenge wiederum kreuzungsfrei ist. Im entsprechenden Transformationsgraphen wird jedes Paar von kompatiblen Konfigurationen durch eine Kante verbunden. Da die Transformationsregel eine gewisse Ähnlichkeit der Konfigurationen beschreibt, interessieren wir uns für den Durchmesser des Transformationsgraphen. Wir konstruieren Mengen von n Punkten, sodass der Durchmesser der dazugehörigen Transformationsgraphen von kreuzungsfreien Spannbäumen und perfekten Paarungen jeweils $\Omega(\log n/\log\log n)$ ist. Die bekannte obere Schranke ist in beiden Fällen $O(\log n)$. Messen wir den Durchmesser des Transformationsgraphen der kreuzungsfreien Spannbäume in Abhängigkeit von der Anzahl k konvexer Schichten ist unser Ergebnis scharf, d.h. der Durchmesser ist $\Omega(\log k)$, was bestmöglich ist.

Abschliessend beschäftigen wir uns mit algorithmischen Aspekten, die das Abzählen der kreuzungsfreien geometrischen Graphen betreffen. Wir beweisen, dass die Berechnung der Anzahl solcher kreuzungsfreien Graphen exponentiell schneller möglich ist als ihr Aufzählen benötigt. Dies ist insbesondere deswegen interessant, weil für andere prominente Graphklassen, wie zum Beispiel Triangulierungen, Spannbäume oder perfekte Paarungen, keine analogen Aussagen bekannt sind. Als weiteres Resultat erhalten wir eine untere Schranke für die Anzahl der planaren Graphen in Abhängigkeit von der Anzahl Triangulierungen: es gibt mindestens $\sqrt{8}^{n-1}$ mal mehr kreuzungsfreie Graphen als Triangulierungen auf jeder Menge von n Punkten in allgemeiner Lage. Obere Schranken für die Anzahl planarer Graphen werden meist wie folgt hergeleitet: man schätzt die Zahl der Triangulierungen von oben ab und zählt in jeder Triangulierung sämtliche Untergraphen. Für degenerierte Punktmengen kann diese Abschätzung bestmöglich sein, allerdings verbessern wir die Schranke für Punktmengen in allgemeiner Lage. Dadurch erhalten wir auch die zurzeit beste obere Schranke $O(343.106^n)$ für die Anzahl kreuzungsfreier geometrischer Graphen auf einer Menge von n Punkten.

Geometry

They say who play at blindman's buff
And strive to fathom space
That a straight line drawn long enough
Regains its starting place
And that two lines laid parallel
Which neither stop nor swerve
At last will meet, for, strange to tell,
Space throws them both a curve.

Such guesswork lets my hopes abide,
For though today you spurn
My heart and cast me from your side
One day I shall return
And though at present we may go
Our lonely ways, a tether
Shall bind our paths till time be through
And we two come together.

X. J. Kennedy

0 Introduction

0.1 Motivation

Every so often scientists are granted to catch a brief glimpse at paragraphs or, in case someone is very lucky, even at a page of The Book. According to the famous Hungarian mathematician Paul Erdős, it contains the perfect proofs for all mathematical statements, however, access to The Book is very limited.

Particularly fascinating to the author of these lines are the beautiful identities, and of course the elegant proofs thereof, that were found by brilliant minds in certain chapters of The Book that are entitled combinatorics, graph theory and geometry (granted that there is a division into chapters).

Besides the esthetically pleasing nature of many of these observations and results, their corresponding proofs from The Book often bring about new insights and may allow for methods to tackle and solve new families of problems.

In our studies we came across some known results whose derivations probably borrow from The Book; at least their statements led to such belief. We will mention a few such examples in the course of our description and hope that the esteemed reader will share our passion for their beauty.

0.2 The thesis in a nutshell

We consider graphs whose vertices are given by finite point sets in the plane and whose edges connect corresponding endpoints without passing through any other point. In particular we focus on *geometric graphs* whose edges are drawn as straight-line segments. Such a geometric graph is crossing-free if no pair of its edges shares any point other than a common endpoint. The drawing in Figure 0.1(a) is a crossing-free geometric graph; The embedding in (b) exhibits a point common to the interior of two edges and is thus not crossing-free; Finally, the graph in (c) is not geometric.

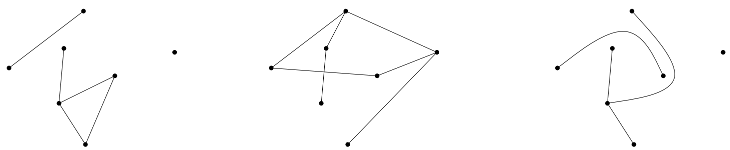

(a) Crossing-free geometric graph (b) Geometric graph (c) Graph embedded in the plane

Figure 0.1: Drawings of graphs in the plane

Throughout the thesis we treat crossing-free geometric graphs even if not stated explicitly. Moreover, to avoid trivialities for the most part general position of the underlying set of points is assumed. Unless stated otherwise this means that no three points are collinear.

In order to avoid confusion when providing estimates for cardinalities we briefly explain our use of the terms *upper* and *lower bound* in this context. For instance let us consider counting the total number of crossing-free geometric graphs that a fixed set of n points allows for.

On the one hand, by definition an upper bound for this quantity cannot be exceeded by any particular choice of n points in the plane. As a measure of quality for such an estimate, one usually exhibits a certain set of n points and determines the number of graphs on this set. At the same time this construction serves as a lower bound indicating by how much the upper bound could be improved in case it is not tight.

On the other hand, similarly, a lower bound for the total number of crossing-free geometric graphs is understood to imply that every set of n points has at least that many graphs. We will provide a corresponding upper bound by a concrete choice of a set of n points which constitutes the space for improvements in case the lower bound is not tight.

In the following we outline our contribution and put the findings into perspective by citing known results and referring to related work.

0.2.1 Crossing-free partitions

In the first part of this thesis we are concerned with partitions of a given finite set of points such that the convex hulls of the individual parts are pairwise disjoint. Such an object is called crossing-free partition. We note that a canonical interpretation as a plane graph is immediate: The set of vertices is the underlying point set, and the edges are given by the segments forming the boundaries of the convex hulls of the partition classes. The illustrations in Figure 0.2 show geometric graphs corresponding to partitions of the underlying point set. We note that in Figure 0.2(c) the partition is crossing-free if the individual parts are of size 4, 2, 1, 1 and it is not crossing-free if the parts have sizes 3, 2, 1, 1, 1.

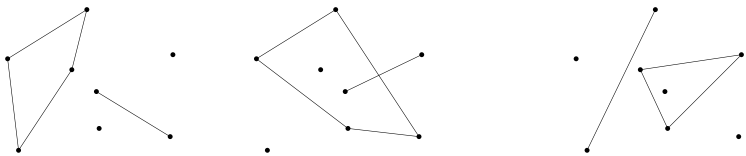

(a) Crossing-free partition (b) Partition which is not crossing-free (c) Crossing-free if and only if parts are of size 4, 2, 1, 1

Figure 0.2: Partitions of points and their interpretation as geometric graphs

Crossing-free partitions of vertex sets defined by convex polygons made their first appearance in the guise of planar rhyme schemes in a note by Becker [18] from the early 1950s. A rhyme scheme is a pattern of lines, referred to by letters of the alphabet, that indicate which lines rhyme. A planar rhyme scheme avoids the pattern *abab* for every subsequence consisting of four lines. Implicitly assumed is a periodic repetition of the complete rhyme scheme. For instance, according to this definition the poem from the quote to this introduction on page 1 constitutes a non-planar scheme.

Becker observed that the number of n-line planar rhyme schemes is the n-th Catalan number. Stanley [77] lists, in an addendum to [76], more than 170 such combinatorial interpretations of the Catalan numbers. Let Γ_n denote the vertex set of a convex polygon on n vertices, and $\mathrm{cfp}(\Gamma_n)$ the number of crossing-free partitions of Γ_n. Then this interpretation precisely means that

$$\mathrm{cfp}(\Gamma_n) = C_n = \frac{1}{n+1}\binom{2n}{n} = \Theta\left(\frac{4^n}{n^{3/2}}\right),$$

where C_n denotes the n-th Catalan number. Moreover, we write $\mathrm{cfp}_k(\Gamma_n)$ for the number of crossing-free partitions of Γ_n into k partition classes.

Their cardinality is known due to Kreweras [49] who, 20 years after Becker's work, was able to infer

$$\mathrm{cfp}_k(\Gamma_n) = \frac{1}{n}\binom{n}{k}\binom{n}{k-1}.$$

This is another notable integer sequence known as the Narayana numbers. Several derivations of these counting results have later been given most of them combinatorial, see for instance Edelman [29], Prodinger [62] or Liaw et al. [53]. Flajolet and Noy [31] obtain the identities invoking the machinery of generating functions.

The study of crossing-free partitions has been fruitful to many disciplines in combinatorics with relations even as far as to molecular biology. However, a comprehensive discussion exceeds the scope of this thesis. We just mention a survey by Simion [73] covering many topics that illustrate the broad variety of impact that crossing-free partitions had since their original treatment by Becker and Kreweras.

Our motivation to investigate crossing-free partitions stems from numerous observations that convex position of a point set minimizes the number of certain crossing-free configurations. As example we briefly mention an argument regarding Hamiltonian cycles. It is not hard to convince oneself that Γ_n allows for one such cycle only, whereas the shortest Hamiltonian cycle on any set of points is always crossing-free. Thus, among all sets of n points in the plane, convex position minimizes the number of crossing-free Hamiltonian cycles.

García et al. [33] proved a corresponding statement for perfect matchings and spanning trees, and recently Aichholzer et al. [8] extended these results by showing that Γ_n also minimizes the number of crossing-free spanning paths, pseudo-triangulations, pointed pseudo-triangulations, forests, connected graphs, and all plane graphs.

However, it is a well-known result due to Hurtado and Noy [39] that triangulations are a prominent counterexample to this pattern. The authors suggest a construction of a set of n points, in the literature referred to as double-circle, which asymptotically allows for $\Theta^*(\sqrt{12}^n)$ triangulations (polynomial factors in n are neglected in the Θ^* notation). This compares to the number of triangulations on Γ_n which is given by the Catalan numbers, and thus of order $\Theta^*(4^n)$. We recall that the asymptotic growth of $\mathrm{cfp}(\Gamma_n)$ is determined by the Catalan numbers as well.

Despite the vastly treated notion of crossing-free partitions of Γ_n only few results are known for arbitrary point sets in general position. Sharir and Welzl [71] show an upper bound of $O(12.24^n)$ for the quantity under consideration, given any set of n points in the plane, and they analyze the so-called double-chain, introduced in [33], to prove the existence of a point set with $\Omega(5.23^n)$ crossing-free partitions. It remains an open problem whether Γ_n minimizes the total number of crossing-free partitions over all point sets in general position but we conjecture an affirmative answer.

On the other end, in our initial result of Chapter 1 we establish that for no $n \geq 5$ does the set Γ_n maximize the number of crossing-free partitions. We exhibit a concrete construction permitting strictly more partitions.

As a first step towards resolving the conjecture we will derive that n points in convex position attain the minimum number of crossing-free partitions into k classes, for certain values of k. In fact, we also conjecture this stronger statement to be true for all k, and give proofs for $k \in \{1, 2, 3, n-3, n-2, n-1, n\}$. Both claims have computationally been verified for sets of at most nine points in general position with the help of the order type database developed by Aichholzer et al. [4]. Alon and Onn [13] showed that a set of n points allows for at most $O(n^{6k-12})$ crossing-free partitions into k classes, for k constant, and this bound is tight.

Note that the Narayana numbers representing $\mathrm{cfp}_k(\Gamma_n)$ are symmetric in the sense that $\mathrm{cfp}_k(\Gamma_n) = \mathrm{cfp}_{n-k+1}(\Gamma_n)$, for all $1 \leq k \leq n$. We will see that such an identity does not necessarily hold for arbitrary sets. For fixed $k \geq 3$, we provide a construction for n points which has, by a factor of order $\Omega(n^2)$, more crossing-free partitions into k classes than into $n-k+1$ classes. Even for $k = o(\log n)$ a factor of $\Omega(n)$ for this ratio is achieved. As the main tool in our proof we adapt the notion of halving edges proposed by Lovász [55] and, in the more general setting of k-edges, by Erdős et al. [30] in the early 1970s. We refer to Matoušek [56, Chapter 11] and the survey of Wagner [81] for a detailed background on k-edges and their applications.

In Chapter 2 we study the question of how many crossing-free partitions we need to decompose the complete graph K_n embedded on a set Γ_n in convex position. Figure 0.3 shows such a minimum decomposition of K_5.

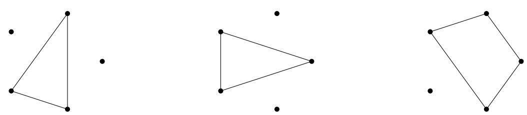

Figure 0.3: Decomposition of K_5 into three crossing-free partitions

Decompositions of graphs are a well-established concept in the area of extremal graph theory. Research in this field was triggered by the Turán problem [78] asking for the maximum number of edges in a graph which does not contain a copy of a given graph H as subgraph. For a deeper introduction to the matter of extremal combinatorics we refer to Jukna [42]. However, we insist on mentioning a result and some open questions concerned with the decomposition of graphs. We note that these problems are stated for abstract graphs without the crossing-free or geometric framework we usually assume in this thesis.

Lovász [54] showed that every graph on n vertices allows for a decomposition into $\lfloor n/2 \rfloor$ paths and cycles. It is, however, not known whether for a decomposition into paths alone at most $\lceil n/2 \rceil$ paths suffice. This is conjectured to be true by Gallai, and the analogue bound for a decomposition into at most $\lfloor n/2 \rfloor$ cycles is conjectured by Hajós.

Ringel [67] proposed an innocent looking problem which also remains unsolved to this day: For any fixed tree T on n vertices K_{2n-1} decomposes into copies of T. Alspach [14] considers decompositions of K_n into cycles of prescribed length. To be more precise, he conjectures that for n odd and $c_1, \ldots, c_k$ natural numbers between 3 and n summing to $\binom{n}{2}$, there is a decomposition of K_n into cycles of lengths $c_1, \ldots, c_k$.

In our setting we attempt to decompose K_n embedded on a convex polygon by means of crossing-free partitions. We show that for the complete graph on Γ_n, and n sufficiently large, at least $n-4$ crossing-free partitions are necessary for that purpose. Conversely, a construction solely using maximal crossing-free matchings on Γ_n results in a decomposition of K_n with exactly n partitions providing an almost tight upper bound.

Our results presented in Chapter 1 are joint work with Emo Welzl [65], and the results in Chapter 2 are obtained jointly with Sonja Čukić, Michael Hoffmann, and Tibor Szabó [24].

0.2.2 Transformation graphs

In Part II we introduce the abstract framework of transformation graphs which will accompany us throughout the remainder of the thesis. For a finite set P of n points in the plane we consider a graph whose vertices $\mathcal{F} = \mathcal{F}(P)$ are certain crossing-free configurations on P, and its edges connect configurations that may be obtained one from the other by predefined rules of transformation. This transformation graph is denoted $\mathcal{T}_{\mathcal{F}}(P)$. The transformation rules bring about a notion of similarity which motivates to investigate the diameter of $\mathcal{T}_{\mathcal{F}}(P)$.

For example, a well-studied object of this type is the graph $\mathcal{T}_{\mathrm{tr}}(P)$ defined on the set of triangulations on P, where two of its elements are adjacent if one is obtained from the other by flipping an edge in the respective triangulation, see Figure 0.4. In his seminal works Lawson showed that, for any set P of n points, $\mathcal{T}_{\mathrm{tr}}(P)$ is connected [50] and has diameter $O(n^2)$ [51]. Hurtado et al. [40] proved this upper bound to be asymptotically tight in the worst case, but it remains an open problem to find matching constants for the number of edge flips when transforming triangulations.

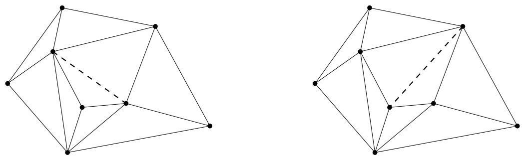

Figure 0.4: Two triangulations adjacent in $\mathcal{T}_{\mathrm{tr}}(P)$

Hanke et al. [34] showed that the distance of two triangulations in $\mathcal{T}_{\mathrm{tr}}(P)$ is at most the number of crossings between the edges of the two triangulations. This also implies the quadratic upper bound on the diameter since every triangulation has less than $3n$ edges and therefore the maximum number of intersections is at most $9n^2$. In addition, the result favors the intuition about short transformation sequences in case the triangulations are very similar, i.e., there are only few intersections between the edge sets, as opposed to flipping both triangulations to a third canonical one (for instance to the Delaunay triangulation as in [51]).

Let us point out that such flips in triangulations play a fundamental role for the enumeration of different configurations of crossing-free geometric graphs, see for instance [15, 20]. For a general treatment covering flips in geometric and abstract graphs we refer to the recent survey by Bose and Hurtado [21].

We compare these results to the case of a point set Γ_n in convex position. For $n > 12$, a tight linear upper bound of $2n - 10$ for the diameter of $\mathcal{T}_{\mathrm{tr}}(\Gamma_n)$ is known due to Sleator et al. [74]. They related this diameter to the minimum number of rotations needed to convert two binary trees on $n-2$ nodes into one another. In fact, they hereby improved over a previous result of Culik and Wood [25] who were the first to consider the number of rotations for converting trees on n nodes. Moreover, Sleator et al. showed their bound to be tight for an infinite set of values of n.

Now, we draw our attention to crossing-free spanning trees and perfect matchings which we will treat in this thesis. Avis and Fukuda [15] define the graph $\mathcal{T}^1_{\mathrm{st}}(P)$ on the crossing-free spanning trees of P where two trees are adjacent if their symmetric difference is a path of length 2 starting at the leftmost point of P. They show that this transformation graph is connected for any P and has diameter at most $2n-4$.

Hernando et al. [37] consider the transformation graph $\mathcal{T}_{\text{st}}^2(\Gamma_n)$ on points in convex position where two trees are adjacent if their symmetric difference is of size 2. They prove that $\mathcal{T}_{\text{st}}^2(\Gamma_n)$ is Hamiltonian and its connectivity is equal to the minimum vertex degree. Moreover, the authors give a lower bound of $3n/2-5$ on the diameter of $\mathcal{T}_{\text{st}}^2(\Gamma_n)$.

Another transformation graph of spanning trees on P, suggested by Aichholzer et al. [3], results from the edge-slide operation on a tree. Two trees T and T' are adjacent in $\mathcal{T}_{\text{st}}^3(P)$ if there is an edge of T such that keeping one of its endpoints fixed and sliding the other endpoint along a respective adjacent edge yields T', see Figure 0.5. It is shown that $\mathcal{T}_{\text{st}}^3(P)$ is connected, which has the important consequence that spanning trees can be transformed into each other applying local, constant-size operations only. A polynomial upper bound of $O(n^2)$ on the corresponding diameter was recently given by Aichholzer and Reinhardt [11].

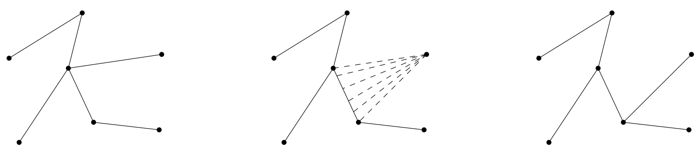

Figure 0.5: The edge slide transformation on spanning trees

Two crossing-free graphs on P are said to be compatible if the union of their edge sets is again without crossing. We note that convex position of the underlying point set Γ_n is of minor interest for the transformation of compatible trees or matchings, as the diameter of the corresponding transformation graph is 2 if $n \geq 4$, and 1 otherwise. Indeed, in both cases there is a universal crossing-free representative consisting of edges from the boundary of the convex hull which is adjacent to any other configuration, see Figure 0.6.

The authors of [3] also consider a directed transformation graph $\mathcal{T}_{\text{st}}^4(P)$ on spanning trees: A tree connects via a directed edge to the compatible tree of minimum Euclidean length. In this setting $\mathcal{T}_{\text{st}}^4(P)$ is shown to be a rooted tree, with the Euclidean minimum spanning tree of P as root, and any tree has distance at most $O(\log n)$ from the root.

This bound was refined by Aichholzer et al. [2] for the undirected transformation graph $\mathcal{T}_{\text{st}}^5(P)$ of compatible spanning trees by proving an upper bound of $O(\log k)$ for the diameter of $\mathcal{T}_{\text{st}}^5(P)$, where k denotes the number

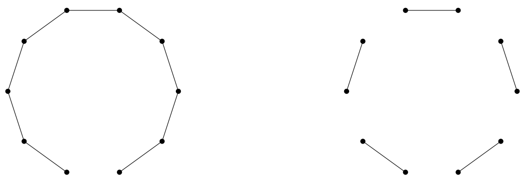

Figure 0.6: Universal spanning tree and perfect matching on Γ_n

of convex layers of P. Their work was also motivated by showing that an upper bound of d on the diameter of compatible spanning trees implies an upper bound of $O(nd)$ on the diameter of the flip graph of pseudo-triangulations of P. The authors conjectured the diameter of $\mathcal{T}_{\text{st}}^5(P)$ to be sub-logarithmic and mentioned that no pair of spanning trees is known whose distance is more than a constant.

We partially close this gap in Chapter 3 by suggesting a construction for point sets P of cardinality n with k convex layers that allow for two spanning trees whose distance in $\mathcal{T}_{\text{st}}^5(P)$, in terms of the number of convex layers, achieves a tight lower bound of $\Omega(\log k)$ and an almost tight bound of $\Omega(\log n / \log\log n)$, in terms of the order of P.

Houle et al. [38] considered the transformation of compatible perfect matchings. They showed that for any set P of an even number n of points the corresponding graph $\mathcal{T}_{\text{pm}}(P)$ is connected and has diameter at most $n-2$. In a very recent work Aichholzer et al. [6] improved the bound to $O(\log n)$. In Chapter 4 we complement this upper bound by a construction yielding a sub-logarithmic lower bound of $\Omega(\log n / \log\log n)$.

Hernando et al. [36] analyze the transformation graph for perfect matchings on Γ_n where the adjacency relation is given by a symmetric difference of exactly two edges. They show that the graph is bipartite, of diameter $n-2$, and Hamiltonian if $n/2$ is even, and does not contain a Hamiltonian path for odd $n/2$.

The results on the transformation of spanning trees are joint work with Kevin Buchin, Uli Wagner and Takeaki Uno [22].

0.2.3 Counting crossing-free geometric graphs

The final part of this thesis is devoted to counting the number of crossing-free geometric graphs on a set of n points in the plane. Central to both Chapters 5 and 6 is the fact that this quantity never exceeds a fixed exponential in n. Ajtai et al. [12] were the first in 1982 to establish this result with 10^{13} as base of the exponential. Hereby, they answered a question raised by Newborn and Moser [61] asking for an upper bound on the number of crossing-free spanning cycles in a set of n points. This original problem led to the discovery of the celebrated *Crossing Lemma* by Ajtai et al. [12], which asserts a strong lower bound for the number of crossings in a geometric graph with many edges. Its first of numerous applications later on was the exponential bound on the number of plane graphs.

On the other end, it is clear that any set of n distinct points allows for exponentially many crossing-free graphs. Indeed, the shortest spanning tree is always crossing-free and, hence, so are its 2^{n-1} subgraphs. We note that this is even true for points not in general position. Likewise, it was commonly believed that a point set in general position has exponentially many triangulations, however, a rigorous argument settling this question and a lower bound of 2.012^n was only lately given by Aichholzer et al. [9] in 2001. (The same authors also mention a previous result of Galtier et al. [32], albeit published later than [9], which implies a bound of $\Omega(1.124^n)$.)

The most recent results concerning lower bounds on the number of triangulations in any set of n points are the general estimate $\Omega(2.338^n)$, due to Aichholzer et al. [10], and for the particular case of k points on the boundary of the convex hull McCabe and Seidel [58] show a lower bound of $\Omega\left((\frac{30}{11})^k(\frac{11}{5})^{(n-k)}\right)$. For k fixed, the latter bound yields $\Omega(2.2^n)$ but can be improved to $\Omega(2.63^n)$.

By applying the reverse search technique of Avis and Fukuda [15] it is possible to enumerate the set of triangulations on a set P of n points in time at most a polynomial factor in n times the number of triangulations. Our contribution in Chapter 5 is that we show how to count the number of crossing-free geometric graphs on a given point set exponentially faster than enumerating them. More precisely, given a set P of n points in general position we may compute $\mathsf{pg}(P)$, the number of crossing-free geometric graphs on P, in time at most $\frac{\mathrm{poly}(n)}{\sqrt{8}^n} \cdot \mathsf{pg}(P)$. It is worth mentioning that no similar statements are known for other prominent graph classes like triangulations, crossing-free spanning trees or perfect matchings.

To achieve the exponential speed-up we assign every triangulation, on average, to at least $\sqrt{8}^{n-1}$ crossing-free geometric graphs. For point sets with triangular convex hulls we are able to improve the base of the exponential from $\sqrt{8} \approx 2.828$ to 3.347. As main ingredient for the improvement we show that there is a constant $\alpha > 0$ such that a triangulation on P chosen uniformly at random contains, in expectation, at least n/α non-flippable edges. The latter result is of interest in its own right, and the best value for α we obtain is $37/18$.

We mentioned the upper bound of 10^{13n} by Ajtai et al. [12] on the number of crossing-free graphs on a set of n points. Further stimulus in this research area [75, 70, 26, 69] was mainly motivated by progressively deriving better bounds, 173000^n, 4854.52^n, 276.76^n, 59^n for the number $\mathsf{tr}(P)$ of triangulations on n points, where the currently best known upper bound stands at 43^n due to Sharir and Welzl [72] from 2006.

Upper bounds on the number of crossing-free graphs are then usually obtained by counting all subsets of edges in each triangulation. Let M denote the number of edges in any triangulation, then this amounts to multiplying the bound on the triangulations by 2^M.

Lastly, in Chapter 6 we derive the first non-trivial upper bound of the form $2^{\gamma \cdot M} \cdot \mathsf{tr}(P)$, with $\gamma < 1$, for the number of crossing-free geometric graphs on a set P in general position, and we deduce a bound of $O(343.106^n)$ for $\mathsf{pg}(P)$. The important point to note here is that general position is crucial, as degeneracies may indeed cause a ratio of 2^M for the number of crossing-free graphs versus that of triangulations.

As main component in our derivation we show that there is a constant $\beta \geq \frac{1}{144} > 0$ such that, for any set of at least five points in general position, a crossing-free geometric graph that is chosen uniformly at random contains, in expectation, at least $(\frac{1}{2} + \beta)M$ edges.

The findings in Chapter 5 are joint work with Emo Welzl [66], and the results of Chapter 6 are jointly with Jack Snoeyink and Emo Welzl [64].

0.3 Terminology and Notation

The reader who is experienced in the field of combinatorics will be familiar with most of the notation employed in this thesis. Thence, she may safely skip this section as we briefly review standard mathematical objects and terminology, however, it might still serve as a reference. Specific notions tailored to support the analyses of our results will be introduced gradually and even repeated where deemed necessary.

The set of natural numbers not including zero is denoted by $\mathbb{N}$, the set of integers by $\mathbb{Z}$, and the set of real numbers by $\mathbb{R}$. Given a real number $x \in \mathbb{R}$ we write $\lfloor x \rfloor \in \mathbb{Z}$ for the largest integer smaller than or equal to x, and $\lceil x \rceil \in \mathbb{Z}$ for the smallest integer larger than or equal to x. Given a set S we write $|S|$ for its cardinality, and for $k \in \mathbb{N}$ we denote by $\binom{S}{k}$ the k-element subsets of S. A *partition* of a set S is a collection of non-empty disjoint subsets, called *partition classes*, whose union equals S.

The real 2-dimensional space endowed with the standard inner product will typically be referred to as *Euclidean plane* or simply *plane*. It is denoted by $\mathbb{R}^2$, but keep in mind that the object under consideration is a metric space. With a point $p = (x, y) \in \mathbb{R}^2$ its Euclidean norm $\|p\| = \sqrt{x^2 + y^2}$ is associated. Given the topology induced by this norm, for a point set $P \subseteq \mathbb{R}^2$ in the plane the symbol ∂P denotes its *boundary*, P° the *interior* and $\overline{P}$ the *closure* of P. The *convex hull* of P will be abbreviated by $\text{conv}(P)$. Given a finite set P of points in general position and $Q \subseteq P$ a non-empty subset, by $X(Q) := Q \cap \partial\text{conv}(Q)$ we refer to the extreme points of Q, and the points from P contained inside the convex hull of Q are denoted by $I^P(Q) := P \cap \text{conv}(Q)^\circ$.

A finite point set is said to be in *general position* if no three of its points are collinear. We always assume a given point set to be in general position unless stated otherwise. If a particular context requires an even more restricted assumption we will explicitly mention a corresponding definition.

Let A, B, and C be arbitrary non-empty sets. A function $f : A \to B$ is said to be *injective* if there exists a function $g : B \to A$ satisfying $g(f(a)) = a$, for all $a \in A$, and it is *surjective* if there is $g : B \to A$ such that $f(g(b)) = b$, for all $b \in B$. If f is both injective and surjective we call it *bijective*. Let $f : A \to B$ and $g : B \to C$ be functions then we write $g \circ f : A \to C$ for the function defined by $g \circ f(a) := g(f(a))$, for all $a \in A$.

A real-valued function f defined on some real interval A is called *convex* if $f\big((1-t)x + ty\big) \leq (1-t)f(x) + tf(y)$, for all $t \in [0, 1]$ and $x, y \in A$.

We employ the common Landau notation to describe the limiting behavior of a function as its argument tends to infinity. Let $f, g : \mathbb{N} \to \mathbb{R}$ be two functions taking non-negative values only, then we write $f(n) = O(g(n))$, or simply $f = O(g)$, and say $f(n)$ *is of order at most* $g(n)$ if

$$\exists c > 0 \; \exists n_0 \in \mathbb{N} \; \forall n \geq n_0 \; : \; f(n) \leq c \cdot g(n).$$

With slight abuse of notation we also want to allow the use of expressions like $O(g) = O(h)$, if for all function satisfying $f = O(g)$ also $f = O(h)$ holds. We observe here that the symbol "=" is not symmetrically used in this context. Furthermore, we introduce the notion $f = \Omega(g)$ which is equivalent to saying $g = O(f)$, and $f = \Theta(g)$ means both $f = O(g)$ and $f = \Omega(g)$. We often encounter functions growing exponentially fast, in which case we tend to use the notation Θ^* instead of Θ which neglects polynomial factors and just specifies the dominating exponential term. If $\lim_{n\to\infty} \big(f(n)/g(n)\big) = 0$ we write $f = o(g)$.

Whenever a probability space is understood the *probability* of an event A is denoted $\mathbb{P}[A]$, and the *expected value*, or *expectation* for short, of a random variable X is written as $\mathbb{E}[X]$. Given a predicate A we write $\mathbb{1}_{[A]}$ for the indicator function, i.e., $\mathbb{1}_{[A]} = 1$ if A holds and $\mathbb{1}_{[A]} = 0$ otherwise.

We conclude by introducing the graph theoretic terminology where we mostly follow the notation suggested by West [82]. A *graph* G is an ordered pair consisting of a finite *vertex set* $V(G)$ and a finite *edge set* $E(G)$, whose elements are pairs of vertices. We refer to the vertices of such a pair as the edge's *endpoints*. The endpoints of an edge are also called *neighbors* and said to be *adjacent*, while being *incident* to their connecting edge. If the edges in $E(G)$ are ordered pairs we call G a *directed graph*, otherwise G is *undirected*. For an edge e in an undirected graph between vertices u and v we write $e = \{u, v\}$, or $e = uv$ for short.

For a graph G we write $n(G) = |V(G)|$ for its *order*, and $e(G) = |E(G)|$ for its *size*. A *loop* is an edge whose endpoints are equal. Edges having the same pair of endpoints are termed *multiple edges*. A graph without loops and multiple edges is called *simple*. Unless stated otherwise, graphs are considered simple and undirected throughout this thesis.

A *subgraph* H of G is a graph that satisfies both $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$, and we also say G *contains* H. Moreover, we call H a *spanning subgraph* if $V(H) = V(G)$. A *decomposition* of a graph is a set of subgraphs such that each edge appears in exactly one subgraph of the set.

We allow some operations that may be performed on a graph G. *Deleting an edge* $e \in E(G)$ yields a new graph H with $V(H) = V(G)$ and $E(H) = E(G) \setminus \{e\}$. *Adding an edge* between two non-adjacent vertices of G is defined accordingly. *Deleting a vertex* v in G means removing v from $V(G)$ and all edges incident to v in $E(G)$. Finally, the *contraction of an edge* $e = \{u, v\}$ is defined as the replacement of u and v by a single vertex whose incident edges are all edges, distinct from e, that were incident to u or v. We note that we may hereby obtain multiple edges or even loops.

The number of edges incident to a vertex $v \in V(G)$ is called its *degree* and denotes $\deg(v)$. There is an intrinsic relationship between the degrees of all vertices and the size of a graph, the so-called *Handshaking Lemma*, which reads $\sum_{v \in V(G)} \deg(v) = 2e(G)$.

The *complete graph* on n vertices, denoted K_n, is a simple graph where every pair of vertices is adjacent, hence $e(K_n) = \binom{n}{2}$. A set $M \subseteq E(G)$ of edges with no shared endpoints is called a *matching*. If every vertex is incident to an edge in the matching M, then we call M a *perfect matching*. With slight abuse of notation but to simplify terminology, we sometimes let the term matching refer to a subgraph whose edge set is a matching.

A *path* is a simple graph whose vertices may be ordered such that two vertices are adjacent if and only if they occur consecutively in the ordering.

A *cycle* is a simple graph that is obtained from a path by adding an edge joining the two vertices of degree 1. The *length* of a path, or a cycle respectively, is the number of its edges. We say a graph G is *connected* if for every pair of vertices u and v there is a subgraph of G which is a path such that u and v are of degree 1. The length of the shortest such path is the *distance* between u and v. The *diameter* of G is the maximum distance over all pairs of vertices in $V(G)$. It is infinite if G is not connected.

A simple graph is called *forest* if it does not contain a cycle, a connected forest is a *tree*. We say G is *Hamiltonian* if there is a spanning cycle in G.

Given a finite set P of points in $\mathbb{R}^2$, a *geometric graph* is a simple graph defined on the vertex set P whose edges are straight-line segments connecting the corresponding endpoints. We call a geometric graph *crossing-free* or *plane* if in the embedding on P no pair of its edges shares any point except for, possibly, a common endpoint.

A *face* of a crossing-free graph G is a maximal region of $\mathbb{R}^2$ that does not contain points used in the embedding. Let $f(G)$ denote the number of faces then Euler's polyhedral formula states $n(G) - e(G) + f(G) = 2$.

Every finite plane graph has one unbounded face. A crossing-free geometric graph which is maximal with respect to the number of edges, i.e., no edge can be added without incurring a crossing, is termed *triangulation*. Except for the unbounded face, every face in a triangulation is a triangle.

An angle α such that $0 < \alpha < \pi/2$ is called *convex*, and we say α is *reflex* if $\pi/2 < \alpha < \pi$. Consider a plane embedding of a cycle such that at exactly three vertices there is a convex angle with respect to the bounded face. Then this bounded face is called *pseudo-triangle*.

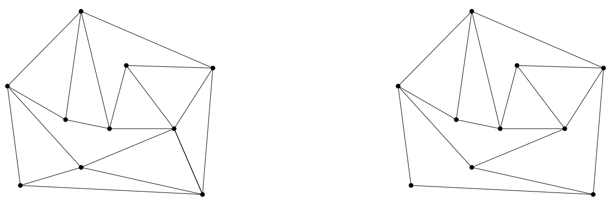

Figure 0.7: Pseudo-triangulation and pointed pseudo-triangulation

A *pseudo-triangulation* is a plane graph where every bounded face is a pseudo-triangle, see the left drawing in Figure 0.7. A *pointed pseudo-triangulation* is a plane graph such that at every vertex there is a reflex angle, and no edge can be added while preserving this property, as in the right picture of Figure 0.7.

Part I

Crossing-Free Partitions

It's that convex position feeling!
The Advertising Slogan Generator

1

Counting Crossing-Free Partitions

A partition of a point set in the plane is called crossing-free if the convex hulls of the individual parts do not intersect. We prove that convex position of a planar set of n points in general position minimizes the number of crossing-free partitions into 1, 2, 3, and $n-3$, $n-2$, $n-1$, n partition classes. Moreover, we show that for no $n \geq 5$ does convex position of the underlying set of n points maximize the total number of crossing-free partitions.

It is known that in convex position the number of crossing-free partitions into k classes equals the number of partitions into $n-k+1$ parts. This does not hold in general, and we mention a construction for point sets with significantly more partitions into few classes than into many.

This is joint work with Emo Welzl [65].

1.1 Preliminaries

Let P be a set of n points in the plane. We assume that P is in *general position*, i.e., no three points are collinear. A partition of P is called *crossing-free* if the convex hulls of the individual parts do not intersect. We observe that one may uniquely identify such a crossing-free partition with a planar straight-line embedded graph on the vertex set P consisting of the edges forming the boundaries of the convex hulls of the partition classes.

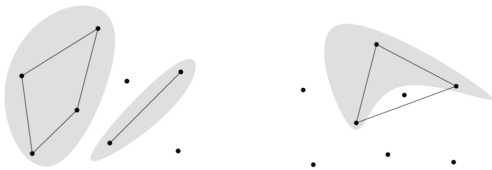

(a) Crossing-free (b) Not crossing-free

Figure 1.1: Partitions of eight points

We denote by $\mathrm{cfp}(P)$ the number of crossing-free partitions of P, and write $\mathrm{cfp}_k(P)$ for the number of crossing-free partitions of P into k classes, where $1 \leq k \leq n$. Moreover, Γ_n denotes a set of n points in convex position, i.e., Γ_n is the vertex set of a convex n-gon.

The notion of crossing-free partitions of Γ_n dates back at least as far as 1952, when Becker [18] in his note on "Planar rhyme schemes" mentioned yet another incarnation of the well-known Catalan numbers. Given our notation his result may be stated in the following way. For any n points in convex position

$$\mathrm{cfp}(\Gamma_n) = C_n = \frac{1}{n+1}\binom{2n}{n} = \Theta\left(\frac{4^n}{n^{3/2}}\right),$$

where C_n denotes the n-th Catalan number. The number of crossing-free partitions into k classes of Γ_n is known due to Kreweras [49] who calculated

$$\mathrm{cfp}_k(\Gamma_n) = \frac{1}{n}\binom{n}{k}\binom{n}{k-1} = \frac{(n-1)!\,n!}{(k-1)!\,k!\cdot(n-k)!\,(n-k+1)!}. \tag{1.1}$$

These numbers also constitute famous integer sequences known as the Narayana numbers which count the number of trees on $n+1$ vertices with exactly k leaves, a result due to Dershowitz and Zaks [27].

Kreweras' original proof reduces to a formal identity shown in his earlier work [48]. Edelman [29] proposed an idea using cyclic permutations and inserting parentheses around partition classes to give a strictly combinatorial derivation of (1.1). Lately, Liaw et al. [53] gave a direct, bijective counting argument based on Kreweras' idea. For an argument involving the concept of Narayana numbers see Prodinger [62], and Klazar [47] with further generalizations. A short derivation of both identities above employing the framework of generating functions can be found in Flajolet and Noy [31].

We will show that Γ_n minimizes $\mathrm{cfp}_k(P)$ for certain values of k if P is in general position, and in fact we conjecture this statement to be true for all k. Note that the term for $\mathrm{cfp}_k(\Gamma_n)$ is symmetric in the sense that $\mathrm{cfp}_k(\Gamma_n) = \mathrm{cfp}_{n-k+1}(\Gamma_n)$, for all $1 \leq k \leq n$. We will see that this is not necessarily the case if the points are not in convex position as, for constant $k \geq 3$, we mention a construction for sets P_k of n points such that

$$\frac{\mathrm{cfp}_k(P_k)}{\mathrm{cfp}_{n-k+1}(P_k)} = \Omega(n^2).$$

In fact, by a result due to Alon and Onn [13] this gap can be made as large as $\Theta(n^{4k-10})$, for partitioning a suitable set of n points in the plane into constant k many classes.

García et al. [33] proved that Γ_n minimizes the number of crossing-free perfect matchings and spanning trees among point sets in general position. Note that Γ_n has $C_{n/2}$ many perfect matchings, a result by Motzkin [60]. Aichholzer et al. [8] extended these results about Γ_n minimizing the number of certain configurations by showing that similar statements also hold for several other graph classes like spanning paths, (pointed) pseudo-triangulations, forests, connected graphs, or all plane graphs. However, it is well-known that triangulations are a prominent counterexample to this pattern, see Hurtado and Noy [39]. It is open whether Γ_n minimizes the total number of crossing-free partitions but we conjecture an affirmative answer. Sharir and Welzl [71] show that $\mathrm{cfp}(P) = O(12.24^n)$, for any set P of n points, and the so-called double-chain, introduced in [33], allows for $\Omega(5.23^n)$ crossing-free partitions.

Proposition 1.1. *For every $n \geq 5$, there is a set S_n of n points in general position such that*

$$\mathrm{cfp}(S_n) > \mathrm{cfp}(\Gamma_n) = C_n.$$

Proof. We define a point set which will prove helpful for several contexts in this thesis. Let S_n denote the *single-chain* formed by n points according to the following construction. For a given circle and a point x outside, let y and z be the two points where the tangents through x touch the circle. Then place $n-3$ points on the circle between y and z, such that the points are contained in the triangle defined by x, y, and z.

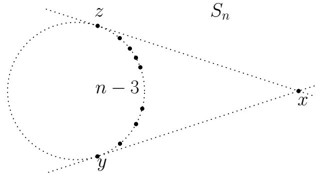

Figure 1.2: Construction of the single-chain S_n

We exhaustively count the number of crossing-free partitions of the hereby obtained point set S_n by separately considering the cases where x belongs to partition classes of size k, for $1 \le k \le n$. By construction, x belongs to a class of size $k \ge 2$ if and only if the other $k-1$ points of this class are consecutive points of the convex set $S_n \setminus \{x\}$. Hence, there are $n-(k-1)$ choices for such a partition class. Furthermore, observe that the remaining $n-k$ points are in convex position, and their individual convex hulls do not intersect the hull of the class containing x. This implies that every crossing-free partition of these $n-k$ points is also crossing-free when additionally considering the partition class containing x. Thus, we find

$$\mathrm{cfp}(S_n) = C_{n-1} + \sum_{k=2}^{n} (n-(k-1)) C_{n-k} =: s_n.$$

Note that $C_{n+1} = 2\frac{2n+1}{n+2} \cdot C_n < 4C_n$ holds for all $n \in \mathbb{N}$. Thus, the claim $s_n > C_n$ is easily established for every $n > 13$, as truncating the sum after the first term yields

$$s_n > C_{n-1} + (n-1)C_{n-2} > C_{n-1} + 12C_{n-2} > C_{n-1} + 3C_{n-1} > C_n.$$

For the remaining values of $5 \le n \le 12$ we simply calculate the exact values of s_n and C_n to find that

$$\begin{array}{ll} s_5 = 43 > 42 = C_5 & s_6 = 141 > 132 = C_6 \\ s_7 = 483 > 429 = C_7 & s_8 = 1704 > 1430 = C_8 \\ s_9 = 6137 > 4862 = C_9 & s_{10} = 22439 > 16796 = C_{10} \\ s_{11} = 82993 > 58786 = C_{11} & s_{12} = 309739 > 208012 = C_{12}, \end{array}$$

which proves the statement. □

1.2 Partitioning into many classes

As a warm-up observe that the number of crossing-free partitions of n points into n and $n-1$ classes does not depend on the order type of the points, as long as they are in general position. A partition into $n-1$ classes corresponds to a plane graph with exactly one edge.

Proposition 1.2. *For a set P of n points in the plane in general position* $\mathrm{cfp}_n(P) = 1$ *and* $\mathrm{cfp}_{n-1}(P) = \binom{n}{2}$.

We need a few more notations in order to deal with many partition classes. Given a finite set P of points in general position and $Q \subseteq P$ a non-empty subset, by $X(Q) := Q \cap \partial\mathrm{conv}(Q)$ we refer to the extreme points of Q, and the points from P contained inside the convex hull of Q are denoted by $I^P(Q) := P \cap \mathrm{conv}(Q)^\circ$. For $k \in \mathbb{N}$, we employ the common notion of $\binom{P}{k}$ for the k-element subsets of P. Given a predicate A we write $\mathbb{1}_{[A]}$ for the indicator function, i.e., $\mathbb{1}_{[A]} = 1$ if A holds and $\mathbb{1}_{[A]} = 0$ otherwise.

Moreover, for $k \in \mathbb{N}$, the number of classes we want to partition the n points of P into, let $n_i \in \mathbb{N}$, for $1 \leq i \leq k$, be such that $\sum_{i=1}^{k} n_i = n$ and $n_1 \geq n_2 \geq \ldots \geq n_k \geq 1$. Given such a sequence an $(n_1, n_2, \ldots, n_k)$-partition of P is a partition of P into k point sets of sizes $n_1, \ldots, n_k$. Accordingly, crossing-free $(n_1, n_2, \ldots, n_k)$-partitions of P are defined.

For instance, $\mathrm{cfp}_n(P)$ equals the number of crossing-free $(1, \ldots, 1)$-partitions of P, and the number of crossing-free $(2, 1, \ldots, 1)$-partitions is $\mathrm{cfp}_{n-1}(P)$. Crossing-free $(2, 2, \ldots, 2)$-partitions, i.e., the number of perfect matchings of P, are minimized in Γ_n due to [33].

We also note that the number of crossing-free $(3, 1, \ldots, 1)$-partitions equals the number of empty triangles in P. This quantity is uniquely maximized in convex position as otherwise, due to Carathéodory's Theorem, some point is contained in a triangle spanned by the remaining points. Concerning lower bounds, Katchalski and Meir [44] showed that there always are at least $\binom{n-1}{2}$ empty triangles on any set of n points in general position. They also proved that there is a constant $c > 0$ such that there exists a point set with at most cn^2 empty triangles. The lower bound on the number of empty triangles in any point set was further improved to $n^2 - O(n \log n)$ by Bárány and Füredi [16] which still is the current state of the art. Constructions of specific point sets with few empty triangles progressively achieved the bounds $2n^2$, for n a power of 2, due to Bárány and Füredi [16], $1.791n^2$ given by Valtr [80], and $1.683n^2$ by Dumitrescu [28]. The currently best known bound of only $1.6196n^2$ empty triangles is due to Bárány and Valtr [17]. It is open whether there are point sets of size n with less than n^2 empty triangles, for n sufficiently large.

Several other well-established notions also relate to such crossing-free $(n_1, n_2, \ldots, n_k)$-partitions. For instance, we note that any choice of four points in P allows for three perfect matchings on this subset. If the corresponding edges do not cross they account for a crossing-free $(2, 2, 1, \ldots, 1)$-partition of P. Otherwise they add to the number of crossings in the embedding of the complete graph K_n on P. The latter quantity is known as the rectilinear crossing number. As we just saw, the total number of crossing-free $(2, 2, 1, \ldots, 1)$-partitions of P and the rectilinear crossing number sum to $3\binom{n}{4}$. Much effort has been put into determining the point configurations minimizing the rectilinear crossing number (thus, maximizing the number of crossing-free $(2, 2, 1, \ldots, 1)$-partitions of P), and estimating its asymptotic behavior. It is known that the minimal rectilinear crossing number is at least $0.37969\binom{n}{4}$ a result due to Aichholzer et al. [7], and at most $0.38056\binom{n}{4}$ as shown by Ábrego and Fernández-Merchant [1].

Concluding let us point out the connection to k-sets. For $k \in \mathbb{N}$, a k-set of P is a subset Q of size k such that there is a line strictly separating Q from its complement. Hence, we find that the crossing-free $(k, n-k)$-partitions exactly correspond to the k-sets of P. Note that these are crossing-free partitions into two classes. We refer to the recent survey article of Wagner [81] for more reading on k-sets and the rectilinear crossing number of point sets.

Both Kreweras [49] and Liaw et al. [53] derived (1.1) by proving that the number of crossing-free $(n_1, n_2, \ldots, n_k)$-partitions of Γ_n is

$$\frac{n(n-1)\cdots(n-k+2)}{\prod_{i\geq 1} a_i!},$$

where $a_i := |\{j \mid n_j = i\}|$ is the number of classes of size i, for $1 \leq i \leq k$.

1.2.1 The case of $n-2$ parts

Before stating our general result for $\mathrm{cfp}_{n-2}(P)$ we show that

$$\mathrm{cfp}_{n-2}(\Gamma_n) = 2\binom{n}{4} + \binom{n}{3}.$$

To see this, besides substituting in identity (1.1), we note that there are only two ways for obtaining a partition of n points into $n-2$ classes, as shown in Figure 1.3 (we will usually refrain from drawing points that belong to partition classes of size 1).

We have to count the number of pairs of crossing-free matching edges, $(2, 2, 1 \ldots, 1)$-partitions, and the (empty) triangles, $(3, 1, \ldots, 1)$-partitions, in Γ_n. Every choice of four points allows for two pairs of matching edges,

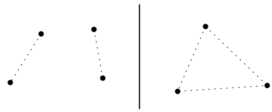

Figure 1.3: Configurations for $n-2$ partition classes

and every set of three points yields a triangle. Hence, the identity from above follows.

In general, however, not every choice of three points in P can result in an empty triangle as mentioned above. Somehow we have to account for the lack of such partitions.

Theorem 1.3. *Let P be a set of n points in the plane in general position. Then*

$$\mathrm{cfp}_{n-2}(P) = 2\binom{n}{4} + \binom{n}{3} + \sum_{Q\in\binom{P}{3}\,:\,|I^P(Q)|\geq 2} \big(|I^P(Q)| - 1\big).$$

In particular we have $\mathrm{cfp}_{n-2}(P) \geq \mathrm{cfp}_{n-2}(\Gamma_n)$.

Proof. We choose $R \in \binom{P}{4}$ and count the number of extreme points $X(R)$. Either $|X(R)| = 4$ and R has exactly two crossing-free perfect matchings, or $|X(R)| = 3$ and R has three perfect matchings. Hence, the number of pairs of crossing-free matching edges in P is

$$\sum_{R\in\binom{P}{4}} 2\cdot \mathbb{1}_{[|X(R)|=4]} + 3\cdot \mathbb{1}_{[|X(R)|=3]}.$$

We doublecount the sets $R \in \binom{P}{4}$ with $|X(R)| = 3$ by iterating over the three extreme points and summing up the number of interior points. Then the term above simplifies to

$$2\binom{n}{4} + \sum_{R\in\binom{P}{4}} \mathbb{1}_{[|X(R)|=3]} = 2\binom{n}{4} + \sum_{Q\in\binom{P}{3}} |I^P(Q)|,$$

and adding the number of empty triangles in P yields that $\mathrm{cfp}_{n-2}(P)$ equals

$$\begin{aligned} & 2\binom{n}{4} + \sum_{Q\in\binom{P}{3}} |I^P(Q)| + \sum_{Q\in\binom{P}{3}} \mathbb{1}_{[I^P(Q)=\emptyset]} \\ = \; & 2\binom{n}{4} + \sum_{Q\in\binom{P}{3}} \Big(1 + \big(|I^P(Q)|-1\big)\cdot \mathbb{1}_{[|I^P(Q)|\geq 2]}\Big) \\ = \; & 2\binom{n}{4} + \binom{n}{3} + \sum_{Q\in\binom{P}{3}} \big(|I^P(Q)|-1\big)\cdot \mathbb{1}_{[|I^P(Q)|\geq 2]}. \end{aligned}$$

Clearly, $\big(|I^P(Q)|-1\big)\cdot \mathbb{1}_{[|I^P(Q)|\geq 2|]} \geq 0$, for all $Q \in \binom{P}{3}$, proving that Γ_n minimizes the number of crossing-free partitions into $n-2$ classes. □

1.2.2 The case of $n-3$ parts

Similar to the case of partitioning P into $n-2$ classes we start by analyzing properties of Γ_n. Note that there are four possibilities to obtain $n-3$ partitions classes, as shown in Figure 1.4. We have to count the number of empty convex quadrilaterals, the triangles containing exactly one point, the empty triangles together with a disjoint edge, and finally the number of triples of crossing-free matching edges. Note that these correspond to $(4,1,\ldots,1)$, $(3,2,1,\ldots,1)$ and $(2,2,2,1,\ldots,1)$-partitions, where a crossing-free $(4,1,\ldots,1)$-partition can be achieved in two ways.

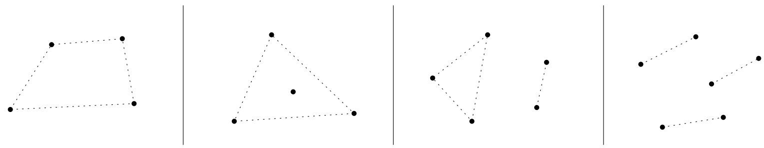

Figure 1.4: All possible configurations to obtain $n-3$ partition classes

Obviously, only three of these configurations may occur if the underlying point set is in convex position. The number of crossing-free $(4,1,\ldots,1)$-partitions is maximized for Γ_n, where the value $\binom{n}{4}$ is attained. Clearly, in general not every choice of four points leads to a crossing-free $(4,1,\ldots,1)$-partition.

A subset of five points from Γ_n gives rise to five distinct crossing-free $(3,2)$-partitions, as shown in Figure 1.5. Since the remaining $n-5$ points from Γ_n do not interfere with such a configuration, i.e., the triangle is always empty, the number of crossing-free $(3,2,1,\ldots,1)$-partitions of Γ_n is $5\binom{n}{5}$. We note that for an arbitrary point set P a crossing-free $(3,2)$-partition of a 5-element subset does not necessarily extend to a crossing-free $(3,2,1,\ldots,1)$-partition of P.

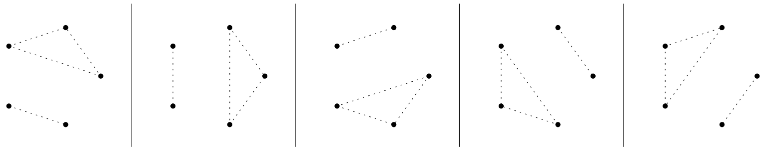

Figure 1.5: Crossing-free $(3,2)$-partitions in convex position

Finally, we count the number of crossing-free $(2,2,2,1,\ldots,1)$-partitions in the convex setup. Then any choice of six points will allow for exactly five such configurations, see Figure 1.6. Clearly, they all extend to crossing-free partitions of Γ_n.

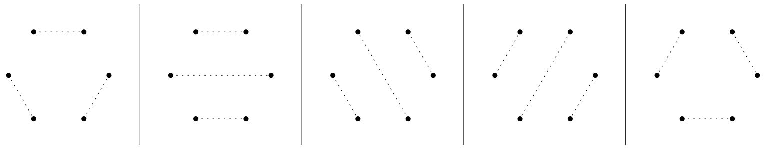

Figure 1.6: Crossing-free $(2,2,2)$-partitions in convex position

Hence, in agreement with evaluating equation (1.1) setting $k = n-3$, we find

$$\mathrm{cfp}_{n-3}(\Gamma_n) = \binom{n}{4} + 5\binom{n}{5} + 5\binom{n}{6}.$$

Now, for an arbitrary set P of n points in general position we will ultimately show $\mathrm{cfp}_{n-3}(P) \geq \mathrm{cfp}_{n-3}(\Gamma_n)$. The proof idea is as follows: For any choice of six points in P there are at least five $(2,2,2)$-partitions which extend to crossing-free $(2,2,2,1,\ldots,1)$-partitions of P. We account for exactly five of them and leave the remaining such configurations for later use. We then consider the crossing-free $(3,2,1,\ldots,1)$-partitions of P and pair them with the previously disregarded $(2,2,2,1,\ldots,1)$-partitions.

More specifically, we choose a subset $Q \in \binom{P}{5}$ of five points. Then $X(Q)$ consists of either five, four or three points, due to the general position assumption. We investigate the $(3,2)$-partitions of Q which do not extend to crossing-free $(3,2,1,\ldots,1)$-partitions of P because of points in $P \setminus Q$

creating a partition class of size at least 4. In order to compensate for such destroyed configurations we count disregarded $(2, 2, 2, 1, \ldots, 1)$-partitions of P instead, where the partition classes of size 2 consist of the five points in Q and a point of $I^P(Q) \setminus Q$. Finally, we will add the crossing-free $(4, 1, \ldots, 1)$-partitions of P.

The following three lemmas prove that there are always enough crossing-free $(2, 2, 2, 1, \ldots, 1)$-partitions to cover for the (possible) lack of crossing-free $(4, 1, \ldots, 1)$ and $(3, 2, 1, \ldots, 1)$-partitions.

Lemma 1.4. *Every subset $Q \in \binom{P}{5}$ in convex position, i.e., $|X(Q)| = 5$, contributes at least* 5 *to* $\mathrm{cfp}_{n-3}(P)$.

Proof. Since $|X(Q)| = 5$ there are exactly five ways to obtain a crossing-free $(3, 2)$-partition of Q, as seen in Figure 1.5. If none of the triangles contains a point from $P \setminus Q$ the claim is true, since every such partition extends to a crossing-free $(3, 2, 1, \ldots, 1)$-partition of P. Otherwise, assume that there is a point in $P \setminus Q$ destroying some triangle.

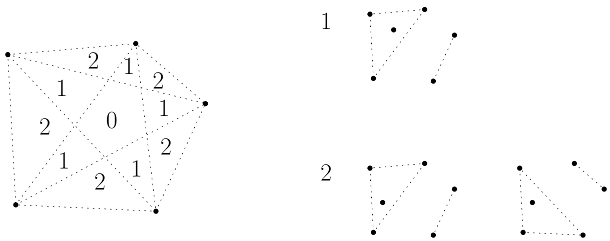

Figure 1.7: Regions for points of $P \setminus Q$, and corresponding destroyed $(3, 2)$-partitions

We associate with each point $x \in \mathrm{conv}(Q)$ its type, which is defined to be the number of 4-element subsets of Q whose convex hulls do not contain x. This results in a subdivision of $\mathrm{conv}(Q)$ into eleven regions of three different types, as depicted in Figure 1.7. Note that a point x of type i, with $0 \leq i \leq 2$, destroys exactly i many $(3, 2)$-partitions of Q. Hence, we may assume that there is a point $x \in I^P(Q)$ of type 1 or 2. We will show that $Q \cup \{x\}$ allows for enough crossing-free $(2, 2, 2)$-partitions to compensate the $(3, 2)$-partitions destroyed by x. Observe that any crossing-free $(2, 2, 2)$-partition of $Q \cup \{x\}$ always extends to a crossing-free $(2, 2, 2, 1, \ldots, 1)$-partition of P.

Consider a partition class of size 2 containing x and one of the points in Q, say y. Then the remaining four points of $Q \setminus \{y\}$ are in convex position, hence allow for exactly two possible crossing-free $(2,2)$-partitions. In the corresponding matchings one of these partitions contains edges only from the boundary of $\mathrm{conv}(Q)$, whereas the other one also contains an edge with points from the interior of $\mathrm{conv}(Q)$.

In the former case, the partition of $Q \setminus \{y\}$ extends to a crossing-free $(2,2,2)$-partition of $Q \cup \{x\}$ for all five choices of $y \in Q$. We consider these five partitions to be the ones we already accounted for in the crossing-free $(2,2,2,1,\ldots,1)$-partitions resulting from every choice of six points in P.

Extending the latter partition of $Q \setminus \{y\}$ to a $(2,2,2)$-partition by adding the edge connection x and y may result in a crossing. However, by definition it follows for all $0 \leq i \leq 2$, that for x of type i there are exactly i points $y \in Q$ such that both $(2,2)$-partitions of $Q \setminus \{y\}$ extend to a crossing-free $(2,2,2)$-partition of $Q \cup \{x\}$, see Figure 1.8.

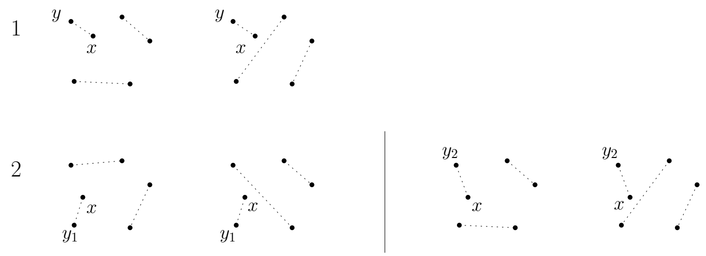

Figure 1.8: Compensating destroyed $(3,2)$-partitions by $(2,2,2)$-partitions

Hence, the number of perfect matchings of $Q \cup \{x\}$, exceeding the five we already accounted for, is exactly the number of $(3,2)$-partitions of Q destroyed by x. If there is more than one point in $I^P(Q)$ of type 1 or 2 then there are even more crossing-free $(2,2,2,1,\ldots,1)$-partitions of P, since for each destroyed $(3,2)$-partition of Q it is enough to consider one point inside the triangle for compensation. This finishes the proof of the lemma. □

We will prove similar results for subsets $Q \in \binom{P}{5}$ with $|X(Q)| = 3, 4$. Observe that in these cases an issue arises which we did not encounter in the previous proof. Namely, there may be distinct 5-element point sets leading to the same set of six points used for compensating destroyed $(3,2,1,\ldots,1)$-partitions when counting $(2,2,2,1,\ldots,1)$-partitions instead. As an example consider Figure 1.9, where the choices of Q are represented by filled dots.

Lemma 1.5. *Every subset $Q \in \binom{P}{5}$ with $|X(Q)| = 4$ contributes at least $5 + \frac{1}{|I^P(Q)|}$ to $\mathrm{cfp}_{n-3}(P)$.*

Figure 1.9: Distinct choices for Q yielding the same six-point configuration

Proof. First let us note that $\left|I^P(Q)\right| \neq 0$ by assumption, hence the term for the contribution of Q is well-defined.

Now, we claim that Q allows for six crossing-free $(3, 2)$-partitions. In order to see this consider the partition class of size 2. Either it corresponds to an edge of the boundary of $\text{conv}(Q)$, for which there are four possibilities, or it connects the point inside the quadrilateral to one of two possible extreme points of Q, see Figure 1.10. The point inside the quadrangle is contained in two triangles with endpoints only in $X(Q)$.

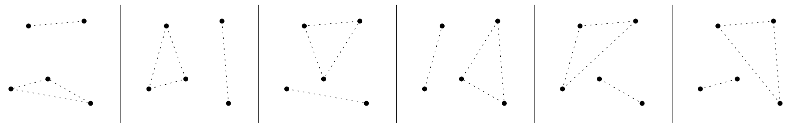

Figure 1.10: Six crossing-free $(3, 2)$-partitions of Q

In order to prove the lemma we show that at least five of these $(3, 2)$-configurations of Q extend to crossing-free $(3, 2, 1, \ldots, 1)$-partitions or may be replaced by crossing-free $(2, 2, 2, 1, \ldots, 1)$-partitions of P, and furthermore, every convex quadrangle in P containing at least one interior point accounts for an additional crossing-free configuration.

Let $y \in Q \setminus X(Q)$ be the point of Q inside the quadrilateral. If no other point from $P \setminus Q$ is inside the quadrangle the statement follows immediately. Otherwise, we make a case analysis distinguishing the relative position of a sixth point $x \in I^P(Q) \setminus Q$ destroying $(3, 2)$-configurations of Q.

In the following figures the initial set Q of five points is represented by filled dots, the point x by a circle and the destroyed triangles by dashed line segments. The diagonals of the quadrangle define four regions where the two interior points x and y may be located. The boundary edges and the diagonals are represented by dotted line segments.

1. Assume that x and y are located in the same region, then only one of the six $(3,2)$-configurations of Q gets destroyed by x. However, there is another choice of five points $Q' := Q \setminus \{y\} \cup \{x\}$ leading to the same six-point configuration, where now y destroys exactly one $(3,2)$-configurations of Q'.

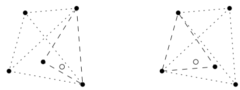

2. If x and y lie in neighboring regions, then x destroys two $(3,2)$-configurations of Q. By symmetry, choosing $Q' := Q \setminus \{y\} \cup \{x\}$ again yields the same six-point configuration, where y destroys two $(3,2)$-configurations of Q'. Hence, in total four such configurations are destroyed.

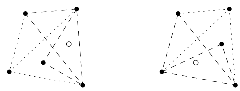

3. If x and y are located in opposite regions, then x destroys three $(3,2)$-configurations of Q, and $Q' := Q \setminus \{y\} \cup \{x\}$ leads to the same six-point configuration, where y destroys three $(3,2)$-configurations of Q'. In total six configurations are destroyed.

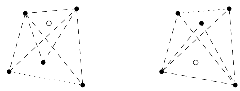

We will now show that in each case the number of crossing-free $(2,2,2)$-partitions of $Q \cup \{x\}$ is large enough to compensate the destroyed $(3,2)$-configurations of Q (and Q'). To this end, we analyze all possible order types of six points with four extreme points.

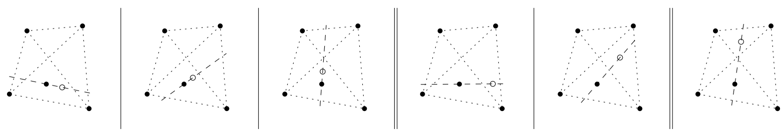

1a 1b 1c 2a 2b 3

Figure 1.11: Relative positions of two points inside a quadrilateral

We need to consider the same three cases as discussed above with additional subcases: Firstly, (1a, 1b and 1c in Figure 1.11) the two interior points lie in the same region defined by the diagonals; in the second case (2a and 2b) the two interior points are located in neighboring regions; in the last case, there is only one order type for the two points to lie in opposite regions.

The following Figure 1.12 demonstrates that the number of crossing-free $(2, 2, 2)$-partitions is eight for the configurations 1a, 1b and 1c.

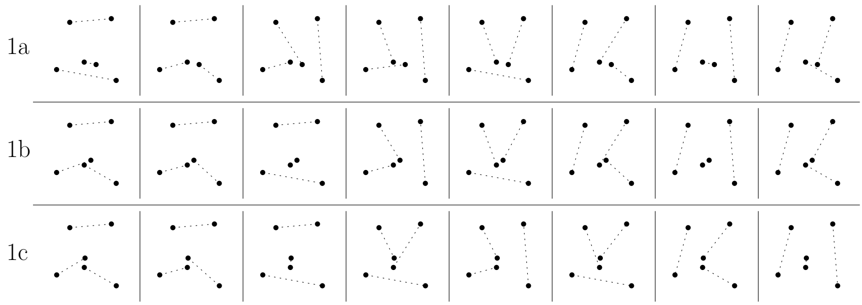

Figure 1.12: Crossing-free $(2, 2, 2)$-partitions for configurations 1a, 1b, 1c

There are nine crossing-free $(2, 2, 2)$-partitions for the configurations 2a and 2b, as shown in Figure 1.13.

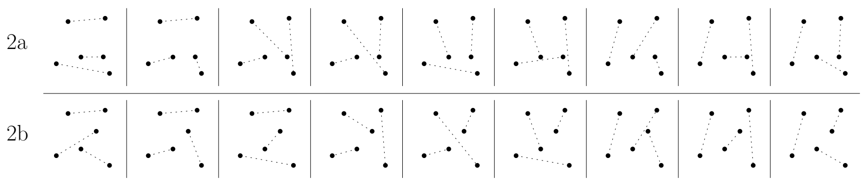

Figure 1.13: Crossing-free $(2, 2, 2)$-partitions for configurations 2a, 2b

Finally, we have ten crossing-free $(2, 2, 2)$-partitions at our disposal for configuration 3, as seen in Figure 1.14.

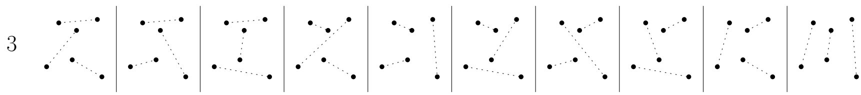

Figure 1.14: Crossing-free $(2, 2, 2)$-partitions for configuration 3

Recall that we already accounted for five of these crossing-free $(2, 2, 2)$-partitions of the six points, and thus the number of crossing-free $(2, 2, 2)$-partitions that remain at our disposal is three, four and five, respectively. By our previous observations we know that in the corresponding cases there are two, four, or six crossing $(3,2,1)$-partitions where we need to count crossing-free $(2,2,2)$-partitions instead.

Hence, in Case 1 and Case 2 all six initial $(3,2)$-partitions can be extended to crossing-free $(3,2,1,\ldots,1)$-partitions or replaced by $(2,2,2)$-partitions for both choices of 5-element subsets Q and Q'. Although in the last case we are lacking one such $(2,2,2)$-partition to restore all six initial $(3,2)$-partitions, we may for each of Q and Q' count five $(3,2)$-partitions together with compensating $(2,2,2)$-partitions of $Q\cup\{x\} = Q'\cup\{y\}$. Moreover, at least one such crossing-free partition remains for the non-empty quadrilateral defined by $X(Q) = X(Q')$ and we get an additional contribution of $1/2$ for both Q and Q'.

If there is more than one point involved in destroying the triangle of a $(3,2)$-partition of Q, we may choose any of these points and make it responsible for compensating with crossing-free $(2,2,2)$-partitions. For each additional point inside the triangle we get even more $(2,2,2)$-partitions. This proves the claim of the lemma. □

We observe that the statement of Lemma 1.5 is tight as seen for the set of six points in Figure 1.14. There the contribution to $\mathrm{cfp}_{n-3}(P)$ of any choice Q of five points is exactly $5 + \frac{1}{|I^P(Q)|}$.

Lemma 1.6. *Every subset $Q \in \binom{P}{5}$ such that $|X(Q)| = 3$ contributes at least 7 to $\mathrm{cfp}_{n-3}(P)$.*

Proof. First observe that the relative position of five points whose convex hull has three extreme points is unique. Hence, we will in the following assume that $Q = \{a,b,c,\ell,r\}$ and its three extreme points are labeled a, b, c and the interior points are ℓ and r such that the segments ar and $b\ell$ cross, see Figure 1.15.

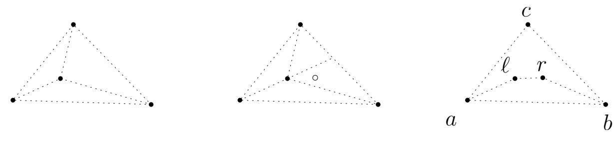

Figure 1.15: Five points whose convex hull has three extreme points

Note that Q allows for seven crossing-free $(3,2)$-partitions, as shown in Figure 1.16. We argue that each of them either extends to a crossing-free $(3,2,1,\ldots,1)$-partition of P, or may be replaced by a crossing-free $(2,2,2,1,\ldots,1)$-partition if points in $P \setminus Q$ destroy the empty triangle.

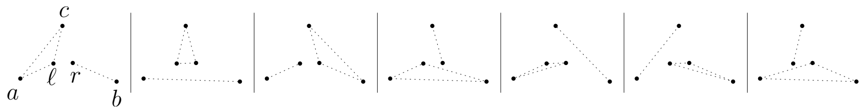

Figure 1.16: Seven crossing-free $(3, 2)$-partitions of Q

If $I^P(Q) \setminus Q$ is empty we are done, thus let $x \in I^P(Q) \setminus Q$ be a point inside the triangle abc. We note that only in the first three partitions of Figure 1.16 the point c belongs to the class of size 3. Hence, if x lies in the quadrilateral $abr\ell$ there are two $(3, 2)$-partitions of Q where x destroys the triangle, and otherwise only one $(3, 2)$-partition of Q gets destroyed by x.

Furthermore, when considering how many times the same six-point configuration for counting crossing-free $(2, 2, 2, 1, \ldots, 1)$-partition may occur we find that, if x lies in the quadrangle $abr\ell$, there are at most six such possibilities, and otherwise we have at most five cases. This is because any choice of such an initial set Q' has to consist of a, b, c and two of the interior points. There are three possibilities to choose two points from $\{\ell, r, x\}$, and in the worst case the third interior point destroys two of the seven $(3, 2)$-partitions, as argued above.

Let us analyze the number of crossing-free $(2, 2, 2)$-partitions of $Q \cup \{x\}$ depending on the position of x.

Suppose that x is not located in the quadrilateral $abr\ell$, then x lies in exactly one of triangles $a\ell c$, ℓrc, or rbc. As it turns out we do not have to further distinguish these cases.

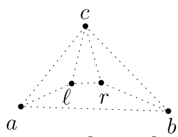

Assume that x belongs to the same partition class as c then the remaining four points are in convex position and allow for two crossing-free $(2, 2)$-partitions, which will not cross the edge between x and c. If x is in the same class as one of the other two points of the triangle, then all three $(2, 2)$-partitions of the remaining four points extend to crossing-free $(2, 2, 2)$-partitions of $Q \cup \{x\}$. Finally, for each of the two remaining points in Q that may lie in the same partition class as x there is at least one crossing-free $(2, 2, 2)$-partitions of $Q \cup \{x\}$. Thus, if x is not inside the quadrangle $abr\ell$, we have at least $2 + 2 \cdot 3 + 2 \cdot 1 = 10$ crossing-free $(2, 2, 2)$-partitions in total.

Now suppose that x lies in the quadrilateral $abr\ell$. Again, the following estimate is independent of the exact position of x.

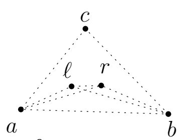

There are two possibilities for x to belong to a class of size 2 together with a point of the quadrilateral, such that this class and all three $(2,2)$-partitions of the remaining four points is still crossing-free. A partition class containing x and one of the other two points of the quadrilateral can be extended in two ways to crossing-free $(2,2,2)$-partitions of $Q \cup \{x\}$. Finally, if x belongs to the same class as c there is at least one crossing-free extension to a $(2,2,2)$-partition. Hence, if x lies inside the quadrangle $abr\ell$ we have at least $2 \cdot 3 + 2 \cdot 2 + 1 = 11$ crossing-free $(2,2,2)$-partitions of $Q \cup \{x\}$.

As we already accounted for five crossing-free $(2,2,2)$-partitions of the six points, the number of crossing-free $(2,2,2)$-partitions that remain at our disposal is at least five if x is not in the quadrangle $abr\ell$, and six otherwise. Turning to our previous observations, we know that in the corresponding cases there are five (six, respectively) six-point configurations where we need to count crossing-free $(2,2,2)$-partitions instead of $(3,2,1)$-partitions. Hence, all seven $(3,2)$-partitions of Q contribute to $\mathrm{cfp}_{n-3}(P)$.

If there is more than one point responsible for destroying a $(3,2)$-partition of Q, the contribution to $\mathrm{cfp}_{n-3}(P)$ is even larger than 7 due to the additional $(2,2,2)$-partitions. This concludes the proof of the lemma. □

Theorem 1.7. *Let P be a set of n points in the plane, in general position. Then*

$$\mathrm{cfp}_{n-3}(P) \geq \mathrm{cfp}_{n-3}(\Gamma_n) = 5\binom{n}{6} + 5\binom{n}{5} + \binom{n}{4}.$$

Proof. The choice of any 6-element subset from P allows for five crossing-free $(2,2,2,1,\ldots,1)$-partitions. Combining the results from the previous Lemmas 1.4, 1.5, and 1.6 on the crossing-free $(3,2,1,\ldots,1)$ and compensating $(2,2,2,1,\ldots,1)$-partitions of P we obtain the contribution from the 5-element subsets. Adding the number of the crossing-free $(4,1,\ldots,1)$-partitions of P, i.e., the number of empty quadrilaterals and the triangles containing exactly one point, we find

$$\begin{aligned}
\mathrm{cfp}_{n-3}(P) \quad \geq \quad & \sum_{Q \in \binom{P}{6}} 5 + \sum_{Q \in \binom{P}{5} \,:\, |X(Q)|=5} 5 \\
& + \sum_{Q \in \binom{P}{5} \,:\, |X(Q)|=4} \left(5 + \tfrac{1}{|I^P(Q)|}\right) + \sum_{Q \in \binom{P}{5} \,:\, |X(Q)|=3} 7 \\
& + \sum_{Q \in \binom{P}{4} \,:\, |X(Q)|=4 \,\wedge\, I^P(Q)=\emptyset} 1 \;+ \sum_{Q \in \binom{P}{3} \,:\, |I^P(Q)|=1} 1.
\end{aligned}$$

Recall that, since P is in general position, the following holds

$$\sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=5} 5 \;+\; \sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=4} 5 \;+\; \sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=3} 5 \;=\; 5\binom{n}{5}.$$

Furthermore, we have

$$\sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=4} \frac{1}{|I^P(Q)|} = \sum_{Q\in\binom{P}{4}} \mathbb{1}_{[|X(Q)|=4\,\wedge\, I^P(Q)\neq\emptyset]},$$

and adding the number of empty quadrangles in P we get

$$\sum_{Q\in\binom{P}{4}} \mathbb{1}_{[|X(Q)|=4\,\wedge\, I^P(Q)\neq\emptyset]} + \sum_{Q\in\binom{P}{4}} \mathbb{1}_{[|X(Q)|=4\,\wedge\, I^P(Q)=\emptyset]} = \sum_{Q\in\binom{P}{4}} \mathbb{1}_{[|X(Q)|=4]}.$$

Doublecounting shows that

$$\sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=3} 7 = \sum_{Q\in\binom{P}{5}\,:\,|X(Q)|=3} 5 + \sum_{Q\in\binom{P}{3}} 2\binom{\left|I^P(Q)\right|}{2}\cdot \mathbb{1}_{[|I^P(Q)|\geq 2]}.$$

Observe that

$$\begin{aligned}
&\sum_{Q\in\binom{P}{3}} 2\binom{\left|I^P(Q)\right|}{2}\cdot \mathbb{1}_{[|I^P(Q)|\,\geq\, 2]} = \\
=\;& \sum_{Q\in\binom{P}{3}} \left|I^P(Q)\right|\cdot \underbrace{\left(\left|I^P(Q)\right|-1\right)}_{\geq 1,\text{ as } |I^P(Q)|\,\geq\, 2} \cdot \mathbb{1}_{[|I^P(Q)|\,\geq\, 2]} \\
\geq\;& \sum_{Q\in\binom{P}{3}} \left|I^P(Q)\right|\cdot \mathbb{1}_{[|I^P(Q)|\,\geq\, 2]},
\end{aligned}$$

and hence, adding the number of triangles containing exactly one point,

$$\sum_{Q\in\binom{P}{3}} 2\binom{\left|I^P(Q)\right|}{2}\cdot \mathbb{1}_{[|I^P(Q)|\,\geq\, 2]} + \sum_{Q\in\binom{P}{3}} \mathbb{1}_{[|I^P(Q)|=1]} \;\geq\; \sum_{Q\in\binom{P}{3}} \left|I^P(Q)\right|.$$

Combining these identities and estimates yields

$$\begin{aligned}\mathrm{cfp}_{n-3}(P) &\geq 5\binom{n}{6}+5\binom{n}{5}+\sum_{Q\in\binom{P}{4}}\mathbb{1}_{[|X(Q)|=4]} + \sum_{Q\in\binom{P}{3}}\left|I^P(Q)\right| \\ &= 5\binom{n}{6}+5\binom{n}{5}+\binom{n}{4} \\ &= \mathrm{cfp}_{n-3}(\Gamma_n),\end{aligned}$$

since any choice of four points in P represents either a quadrilateral or a triangle containing the fourth point. □

1.3 Partitioning into few classes

We define a notion similar to halving edges in order to identify a crossing-free partition with certain 2-element subsets of P. Halving edges were first considered by Lovász [55] and have been studied extensively ever since. The benefit compared to the consideration of the previous section is that only a small number of such subsets is needed for describing crossing-free partitions into few classes.

Let A and B be two disjoint convex polygons in the plane. Then there is a line g separating the two sets, i.e., for all $a \in A$ and $b \in B$ the segment with endpoints a and b intersects g. Now, rotate g counter-clockwise until it becomes tangent to both polygons simultaneously, see Figure 1.17.

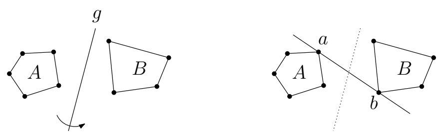

Figure 1.17: Construction of a separating segment

Suppose that no three extreme points of the polygons are collinear then we obtain a unique (extreme) point $a \in A$ and a unique (extreme) point $b \in B$. The segment given by these endpoints a and b is the *separating segment* of the polygons A and B.

Proposition 1.8. *For a set P of n points in the plane in general position* $\mathrm{cfp}_1(P) = 1$ *and* $\mathrm{cfp}_2(P) = \binom{n}{2}$.

Proof. Obviously, $\text{cfp}_1(P) = 1$ holds. To partition P into two crossing-free parts pick a segment defined by two points $p, q \in P$. Then slightly rotate the line through pq clockwise around a point between p and q in order to obtain a separation, as seen in Figure 1.18.

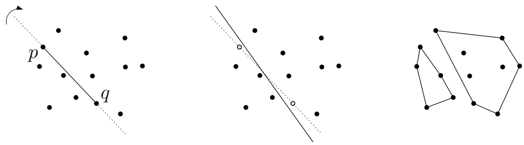

Figure 1.18: Crossing-free partition into 2 classes

Observe that the separating segment for two polygons is unique because of the general position assumption. Thus, for any point set P we have $\text{cfp}_2(P) = \binom{n}{2}$, regardless of the relative position of the points. □

In fact, partitioning a set of points in the plane by a straight line into two classes, or more generally a set in d-dimensional space by a hyperplane, has been considered by Harding [35] who showed that for n points in general position in $\mathbb{R}^d$ there are $\sum_{i=1}^{d} \binom{n-1}{i}$ crossing-free partitions into two classes. In the 2-dimensional plane $\binom{n-1}{1} + \binom{n-1}{2} = \binom{n}{2}$ in agreement with Proposition 1.8.

When partitioning P into more parts it becomes quite handy to attribute colors to the individual classes. Consider the set of surjective maps $\chi : P \to \{1, \ldots, k\}$. Two such k-colorings χ_1, χ_2 are equivalent if there is a permutation π of the k colors such that $\chi_1 = \pi \circ \chi_2$. There is a bijection between the set of equivalence classes of these colorings and the partitions of P into k parts.

If the partition induced by a coloring is crossing-free then for each pair of distinct colors i and j there is a separating segment associated with it. We will indicate the halfplane containing the points of color i (and j, respectively) by drawing i (and j) next to the segment's endpoint in the corresponding halfplane. Moreover, we associate with a crossing-free partition its *separation graph* defined on P whose edges consist of the separating segments of the partition. We observe that this graph need not be crossing-free. A pair of distinct separating segments either does not cross at all, or shares a common endpoint, or intersects in their respective relative interior.

Lemma 1.9. *Given two disjoint separating segments sharing a common color then the segments cannot cross and the coloring of the segments' endpoints receiving the common color is uniquely determined.*

Proof. Without loss of generality let s_{ij} be the segment separating color class i from j, and similarly let $s_{i\ell}$ be the segment separating class i from ℓ. First, observe that s_{ij} and $s_{i\ell}$ cannot intersect in an interior point. Indeed, given the coloring of the endpoints of s_{ij}, see the left-most drawing of Figure 1.19, any coloring of the other segment $s_{i\ell}$ contradicts the definition of separating segments.

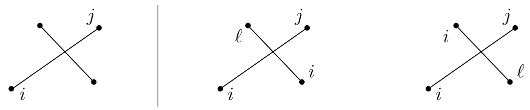

Figure 1.19: Separating segments cannot cross

Since s_{ij} and $s_{i\ell}$ do not intersect, at least one of the segments is completely contained in a halfplane defined by the other segment, which uniquely determines the segments' endpoints of color i, as given in Figure 1.20.

Figure 1.20: Disjoint segments with a common color

This concludes the proof of the lemma. □

When partitioning into $k = 3$ parts any two separating segments share a common color. With Lemma 1.9 we now have the necessary tool to compute the number of crossing-free partitions of a point set into three parts, and we start again by investigating convex position. Observe that there are two ways to partition Γ_n into three classes. Either every class can be separated from both other classes simultaneously by exactly one line, or there is a class that needs two lines, see Figure 1.21.

This difference also arises in the configurations of the corresponding separating segments, see Figure 1.22, where partition classes are drawn as dotted line segments whereas the separating segments are solid.

The crucial observation here is that we only need to know the endpoints of the separating segments and their color in order to reconstruct the underlying partition. Clearly, every three points in Γ_n yield a crossing-free partition into three classes where the corresponding separating segments constitute an empty triangle. Furthermore, every choice of four points

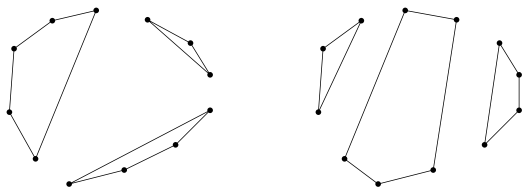

Figure 1.21: Two types of partitioning Γ_n into three parts

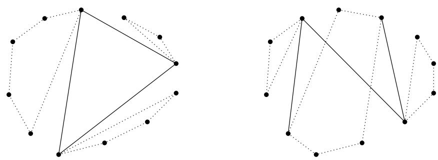

Figure 1.22: The separating segments of the partitions of Γ_n

in Γ_n allows for two different ways of generating a partition into three crossing-free parts, see Figure 1.23. We just showed $\mathrm{cfp}_3(\Gamma_n) = 2\binom{n}{4} + \binom{n}{3}$.

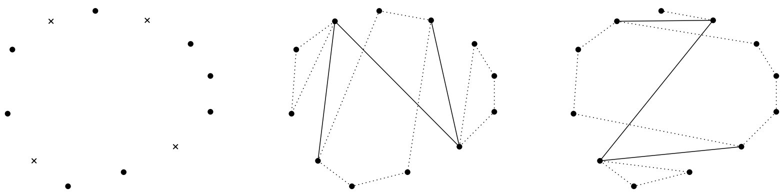

Figure 1.23: Two ways to construct partitions into three classes

If the underlying point set P is not in convex position Carathéodory's Theorem strikes once again, and not every choice of three points leads to an empty triangle. The final goal of our considerations is to establish an injective map from the set of certain separation graphs on P to the crossing-free partitions into three classes. In fact, we will consider the separation graphs arising from the crossing-free partitions of P into $n-2$ parts, where we already know that convex position minimizes their cardinality. Ultimately, we make use of the symmetry $\mathrm{cfp}_{n-2}(\Gamma_n) = \mathrm{cfp}_3(\Gamma_n)$.

We start by showing that every empty triangle accounts for a crossing-free partition of P into three classes.

Lemma 1.10. *There is an injective map from the set of empty triangles* *in P to the crossing-free partitions of P into three partition classes.*

Proof. Let $Q = \{p_1, p_2, p_3\} \subseteq \binom{P}{3}$ with $I^P(Q) = \emptyset$ be an empty triangle in P. Consider the separating segments induced by the edges p_1p_2, p_1p_3, and p_2p_3. Clearly, there is only one way to assign three colors to the three points, up to permutation. The intersection of the halfplanes defined by p_1p_2 and p_1p_3 containing color 1 defines a convex region for the first partition class. Similarly, we find regions where points from the second and third class are located, as seen in Figure 1.24. The only region remaining is the interior of conv(Q) which is empty by assumption.

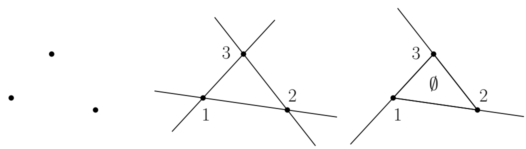

Figure 1.24: Construction from an empty triangle

Hence, we obtain a crossing-free partition of P into three classes. Since the separating segments are unique, two different triangles lead to two distinct partitions. □

Before we turn to the case where a triangle may contain points in its interior we consider the other construction we encountered in the convex setting for three partition classes. To that end we choose four points from P such that their convex hull is a quadrilateral. Note that there are two ways to obtain a pair of disjoint segments from such points and we refer to such a pair as *parallel segments* (even if the lines through the segments intersect).

Lemma 1.11. *There is an injective map from the pairs of parallel segments to the crossing-free partitions of P into three partition classes, and the convex hull of the edges in the separation graph of any such partition is a rectangle.*

Proof. We will construct a crossing-free partition into three classes such that the given segments turn out to be separating segments of the partition. Without loss of generality assume that s_{12} and s_{13} are these segments. By Lemma 1.9 we already know the colors of the segments' endpoints.

We first treat the case where the lines through s_{12} and s_{13} are indeed parallel. The halfplanes defined by these lines divide the plane into three regions whose colorings are defined by the two segments. In this case s_{12} and s_{13} are separating segments of the partition, and we are done.

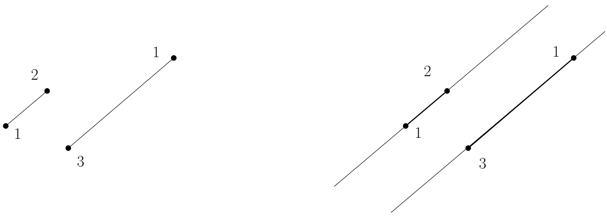

Figure 1.25: Construction from parallel lines

Otherwise assume that the lines through s_{12} and s_{13} intersect in some point. Then the halfplanes defined by the segments divide the plane into four regions one of which did not receive a unique color yet, see the gray region in Figure 1.26. Assigning color 2 or 3 to this region does not violate the coloring induced by the segments s_{12} and s_{13}, however, we choose to color the region with color 3.

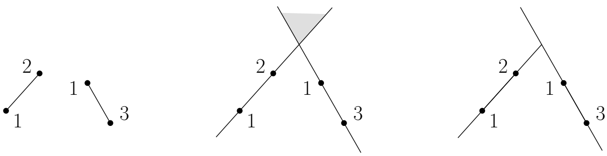

Figure 1.26: Construction from parallel segments

Hence, we found a partition of the plane into three convex regions each of which contains exactly the points of one crossing-free partition class of P, and the given segments s_{12} and s_{13} are separating segments of this partition. Clearly, this partition differs from the ones obtained in Lemma 1.10.

The reason for assigning color 3 to the gray area is that we want the third separating segment s_{23} to pass between the others s_{12} and s_{13}, meaning that each halfplane defined by s_{23} contains exactly one of the other segments. In this way we can ensure that the convex hull of all three segments is a quadrangle.

It remains to show that the map from the pairs of segments to the partitions is injective. As the third segment s_{23} passes between the first two it cannot be parallel to any of them. Now, assume that two pairs of parallel segments map to the same crossing-free partition into three classes.

Since the separating segments of this partition are uniquely defined and by construction exactly two of them are parallel, the pairs have to be the same and the map is injective. □

Finally, we consider configurations that cannot appear in Γ_n which will compensate for both, triangles with points in their interior as well as for pairs of segments whose convex hulls is a triangle. The latter we call *spearing segments.* There are three ways to obtain a pair of disjoint segments from four points with triangular convex hull.

Lemma 1.12. *There is an injective map from the pairs of spearing segments to the crossing-free partitions of P into three partition classes, and the convex hull of the edges in the separation graph of any such partition forms a triangle.*

Proof. Similar to the proof of Lemma 1.11 we want the given segments, without loss of generality s_{12} and s_{13}, to be separating segments of the partition. The colors of the segments' endpoints are given by Lemma 1.9. It remains to determine the partition of the plane which the segments induce. We note that one region defined by the intersection of the corresponding halfplanes does not yet receive a unique color, see Figure 1.27.

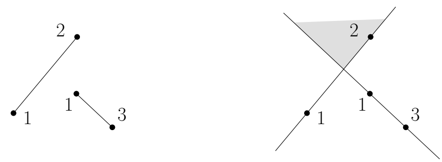

Figure 1.27: Construction from spearing segments

It turns out that in order to obtain an injective map we cannot color the whole gray region with one single color. We need to introduce the third separating segment s_{23} which will determine the mixed coloring of the gray region. For this purpose let y be the endpoint of s_{12} with color 2 and z be the endpoint of s_{13} with color 3. Define x to be the intersection of the lines through the segments s_{12} and s_{13}, and let Δ be the convex hull of the triangle given by x, y, and z, as seen in Figure 1.28.

Rotate the line through x and y counter-clockwise around y and let $p \in P \cap (\Delta^\circ \cup \{z\})$ be the first point of P that gets hit during this rotation. Note that p is well-defined because of general position. We assign color 3 to p and define the third separating segment s_{23} by its endpoints y and p. This results in a partition of the plane into four convex regions where the

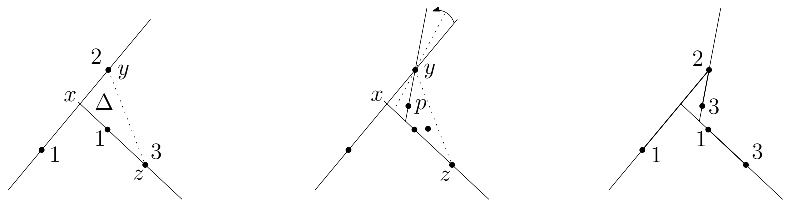

Figure 1.28: Constructing the third separating segment

unbounded areas each contain exactly one partition class and the bounded triangle is empty by construction.

Observe that this partition of the plane into convex regions implies that indeed s_{12}, s_{13}, and s_{23} are the unique separating segments of the underlying crossing-free partition into three classes. Hence, it is clear that the partitions constructed in this way differ from the ones obtained in Lemma 1.11, for here the convex hull of the segments is a triangle. Moreover, they are distinct from the partitions constructed in Lemma 1.10.

It remains to show that two different pairs of spearing segments induce distinct crossing-free partitions. Assume otherwise and note that by construction the convex hull of the two given spearing segments contains the third one. Since the separating segments are unique there is only one choice for two segments such that their convex hull contains the third. Hence, the pairs have to be the same and the map is injective. □

We recall that by considering the separation graphs of the partitions constructed in Lemmas 1.10, 1.11, and 1.12 we find that also the combination of the maps from the three lemmas remains injective.

Theorem 1.13. *Let P be a set of n points in the plane in general position. Then* $\mathrm{cfp}_3(P) \geq \mathrm{cfp}_3(\Gamma_n)$.

Proof. An empty triangle contributes 1 to $\mathrm{cfp}_3(P)$, four points contribute either two pairs of parallel or three pairs of spearing segments. Hence,

$$\begin{aligned}\mathrm{cfp}_3(P) \;&\geq\; \sum_{Q\in\binom{P}{3}\,:\,I^P(Q)=\emptyset} 1 \;+\; \sum_{Q\in\binom{P}{4}\,:\,|X(Q)|=4} 2 \;+\; \sum_{Q\in\binom{P}{4}\,:\,|X(Q)|=3} 3\\ &=\; 2\binom{n}{4} + \binom{n}{3} + \sum_{Q\in\binom{P}{3}} \left(|I^P(Q)|-1\right)\cdot \mathbb{1}_{[|I^P(Q)|\geq 2]}.\end{aligned}$$

The identity follows as in the proof of Theorem 1.3. The last expression is exactly $\mathrm{cfp}_{n-2}(P)$, therefore we get

$$\mathrm{cfp}_3(P) \geq \mathrm{cfp}_{n-2}(P) \geq \mathrm{cfp}_{n-2}(\Gamma_n) = \mathrm{cfp}_3(\Gamma_n). \qquad \square$$

Let us point out that so far we did not specify all possible configurations of separating segments that may occur when partitioning a point set into three crossing-free classes. It will be the result of the remaining part of this chapter that we actually may have left out the largest portion.

1.4 Many partition classes versus few

We start by constructing a set of n points with $\Theta(n^6)$ crossing-free partitions into three classes. This is in contrast to the number of partitions into $n-2$ classes, which always is $\Theta(n^4)$ regardless of the point configuration as long as general position is enforced. Note that for our purpose it is necessary to make use of all three corresponding separating segments of the partition.

In the following we consider a point set that is also known for having a super-linear number of halving edges, a construction and result due to Erdős et al. [30]. For its recursive definition start with the set $P(1)$ of six points, given by a triangle with a smaller similar triangle inscribed and slightly perturbed, see the left illustration in Figure 1.29.

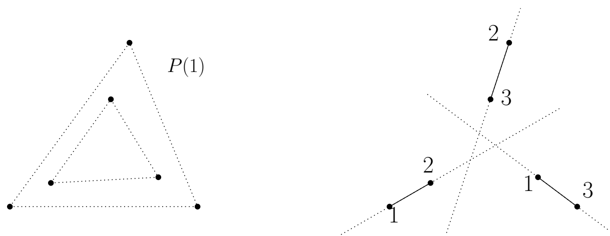

Figure 1.29: Three parallel separating segments

Note that $P(1)$ has a triple of separating segments we did not account for in Theorem 1.13. Here every pair of separating segments is parallel, and by Lemma 1.9 there is only one consistent coloring. If a point set P contains $P(1)$, or an affine copy, as a subset then these three parallel segments additionally contribute to $\mathrm{cfp}_3(P)$ if and only if the bounded region that the lines define is empty. In this case we call the triple *valid*.

Observe that two distinct valid triples of parallel segments induce distinct crossing-free partitions.

Assuming a proper underlying coordinate system, we call a point set *ε-flattened copy of P* if it is obtained by multiplying the y-coordinates of all points in P by ε. Clearly, the slope of any line through two points in an ε-flattened copy of P is ε times the corresponding slope in P. The construction starts with $P(1)$ and, for $\ell \geq 1$, recursively builds $P(\ell+1)$ by arranging three copies of ε_ℓ-flattened point sets of $P(\ell)$ (denoted $P^i_{\varepsilon_\ell}(\ell)$, for $i = 1, 2, 3$) in a tripod, as shown in Figure 1.30.

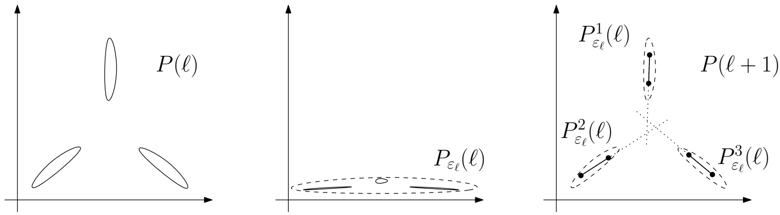

Figure 1.30: Recursive construction of $P(\ell)$

By choosing $\varepsilon_\ell > 0$ small enough we can guarantee that for any choice of six points, two from each $P^i_{\varepsilon_\ell}(\ell)$ with $i = 1, 2, 3$, we obtain a valid triple of parallel segments. Indeed, we can always avoid two points having the same x-coordinate by applying a small perturbation, and then make the slope of any line through two points arbitrarily small by flattening.

For providing a lower bound on $\mathrm{cfp}_3(P(\ell))$ we write $|P(\ell)| = 2\cdot 3^\ell = n(\ell)$ and count the valid triples of parallel segments that can be obtained from $P(\ell)$. This results in

$$\begin{aligned}
\mathrm{cfp}_3(P(\ell)) &\geq \sum_{i=1}^{\ell} 3^{i-1} \cdot \binom{\frac{n(\ell)}{3^i}}{2}^3 = \sum_{i=1}^{\ell} 3^{i-1} \cdot \left(\frac{(\frac{n(\ell)}{3^i})^2}{2}\right)^3 + O(n(\ell)^5) \\
&= \frac{n(\ell)^6}{2^3 \cdot 3} \cdot \sum_{i=1}^{\ell} 3^{-5i} + O(n(\ell)^5) \\
&= \frac{n(\ell)^6}{2^4 \cdot 3 \cdot 11^2} + O(n(\ell)^5) = \frac{n(\ell)^6}{5808} + O(n(\ell)^5).
\end{aligned}$$

For an arbitrary set P of n points the separation graph contains exactly three edges. Hence, there are at most $O(n^6)$ such graphs implying that $\mathrm{cfp}_3(P) = O(n^6)$, and the construction above is asymptotically tight. Conversely, with Theorem 1.13 we find $\mathrm{cfp}_3(P) = \Omega(n^4)$. Recall that $\mathrm{cfp}_k(\Gamma_n) = \mathrm{cfp}_{n-k+1}(\Gamma_n)$ which together with Theorem 1.3 implies that $\mathrm{cfp}_{n-2}(P) = \Theta(n^4)$ for any set P of n points, since $|I^P(Q)| \le n$. This shows that the symmetry which holds for Γ_n is completely lost for arbitrary point sets.

We conclude by generalizing the previous example to crossing-free partitions into k classes, for $k \ge 3$ constant. Since our primary concern here is the asymptotic behavior of the number of crossing-free partitions the following construction is not recursive.

Whereas previously three times $n/3$ points were arranged in a tripod, we now define a set P_k of n points by placing k flattened copies of n/k points in general position in a k-pod. Correspondingly, any choice of $2k$ points, two from each copy, yields a valid k-tuple of parallel segments, and hence contributes to $\mathrm{cfp}_k(P_k)$. With k constant, we obtain the lower bound

$$\mathrm{cfp}_k(P_k) = \Omega\left(\binom{\frac{n}{k}}{2}^k\right) = \Omega\left(\frac{n^{2k}}{2^k k^{2k}}\right) = \Omega\left(n^{2k}\right).$$

On the other hand, in order to bound the number of partitions into $n-k+1$ classes, briefly reflect the derivation of Theorem 1.7. For any set P of n points in general position it holds that $\mathrm{cfp}_{n-3}(P) = \Theta(n^6)$. Indeed, from Figure 1.4 we find that $\mathrm{cfp}_{n-3}(P)$ is upper-bounded by $\binom{n}{6}$ times the maximum number of perfect matchings a set of six points allows for, which is 12, plus terms of lower order $O(n^5)$.

By the same argument, and since k is a constant, there are asymptotically less partitions of P_k into k classes which do not constitute matchings. Thus, $\mathrm{cfp}_{n-k+1}(P_k)$ may be bounded from above by $\binom{n}{2(k-1)}$ times the maximum number of perfect matchings a subset of P_k of size $2(k-1)$ can have. This in turn is at most $10.05^{2(k-1)}$, as shown by Sharir and Welzl [71]. Then again, k is constant, so

$$\mathrm{cfp}_{n-k+1}(P_k) = O\left(\binom{n}{2(k-1)}\right) = O\left(n^{2(k-1)}\right).$$

In fact, this is asymptotically tight since due to [33, 60] the number of perfect matchings in a point set of size $2(k-1)$ is at least C_{k-1}, but a constant. Moreover, partitions into $n-k+1$ classes other than perfect matchings contribute $O(n^{2k-1})$ to $\mathrm{cfp}_{n-k+1}(P_k)$.

If k grows with n, where $k = o(\log n)$, the analog construction and similar estimates show that

$$\frac{\mathrm{cfp}_k(P_k)}{\mathrm{cfp}_{n-k+1}(P_k)} = \Omega(n).$$

Indeed, the lower bound of $\mathrm{cfp}_k(P_k) = \Omega\left(\left(\frac{n}{\frac{k}{2}}\right)^k\right)$ follows in the same way by counting the valid k-tuples of parallel segments, which in fact also holds for all larger k.

In order to derive an upper bound we use another result of [71] which we mentioned in the introduction stating that for set P of n points $\mathrm{cfp}(P) \leq c^n$, for some constant c. Actually, a sufficient condition for our purpose is the fact that the total number of crossing-free graphs on a point set of size n is at most exponential in n, a result which will be discussed in further detail in Chapter 6.

Recalling the notion of $(n_1, n_2, \ldots, n_{n-k+1})$-partitions we may write crossing-free partitions of P_k into $n-k+1$ classes in one of the following forms: $(k, 1, \ldots, 1)$, $(k-1, 2, 1 \ldots, 1)$, …, $(2, \ldots, 2, 1 \ldots, 1)$. Obviously, only points belonging to partition classes of size 2 are incident to edges of the crossing-free partition, and hence we can derive corresponding upper bounds for their cardinality.

There are at most $\binom{n}{k}c^k$ partitions of type $(k, 1, \ldots, 1)$, correspondingly at most $\binom{n}{k+1}c^{k+1}$ of type $(k-1, 2, 1 \ldots, 1)$ and $(k-2, 3, 1 \ldots, 1)$, at most $\binom{n}{k+2}c^{k+2}$ of type $(k-2, 2, 2, 1 \ldots, 1)$, and so on. We already saw that an upper bound for the crossing-free $(2, \ldots, 2, 1 \ldots, 1)$-partitions is given by $\binom{n}{2(k-1)}c^{2(k-1)}$. This yields

$$\mathrm{cfp}_{n-k+1}(P_k) = O\left(\sum_{i=0}^{k-2}\binom{n}{k+i}c^{k+i}\right) = O\left(k \cdot \binom{n}{2(k-1)}c^{2(k-1)}\right),$$

since $k \leq \frac{n}{4}$ and the binomial coefficients are strictly increasing.

Using $(\frac{n}{\ell})^\ell \leq \binom{n}{\ell} \leq (\frac{ne}{\ell})^\ell$ we find

$$\begin{aligned}
\mathrm{cfp}_k(P_k) &= \Omega\left(\frac{n^{2k}}{4^k k^{2k}}\right) \\
\mathrm{cfp}_{n-k+1}(P_k) &= O\left(k\Big(\frac{ne}{2(k-1)}\Big)^{2(k-1)} c^{2(k-1)}\right) = O\left(\frac{n^{2k-2}}{k^{2k}}c'^k\right).
\end{aligned}$$

for some constant $c' \geq c$. With $k = o(\log n)$ the claim follows.

However, for faster growing k, or even linear in n, we can so far not prove a corresponding behavior.

When you wish upon a falling star, your dreams can come true.
Unless it's really a meteorite hurtling to the Earth which will destroy all life.
Then you're pretty much hosed no matter what you wish for.
Unless it's death by meteor.

Wishes
despair, Inc.

2

Decomposing K_n with Crossing-Free Partitions

In this chapter we study how many crossing-free partitions are necessary to decompose the complete graph K_n embedded on a set of n points. Given a point set in general position in the plane a partition is called crossing-free if the convex hulls of the individual partition classes are disjoint.

With every crossing-free partition we associate the following crossing-free geometric graph: The vertices are the given points, and the edges are the convex hull edges of the individual parts. We aim to decompose K_n as an edge-disjoint union of graphs obtained from crossing-free partitions, using the minimal number of such graphs. In particular, we consider an embedding of K_n on Γ_n, a set of n points in convex position.

It is not too hard to see that a decomposition can always be achieved using n crossing-free partitions, more precisely, using crossing-free matchings, i.e., partitions where each part has size 2. We also show that for n sufficiently large, at least $n - 4$ crossing-free partitions are necessary for decomposing K_n, providing an almost tight lower bound for the problem.

The main result of this joint work with Sonja Čukić, Michael Hoffmann, and Tibor Szabó [24] was established during the 4th Gremo Workshop on Open Problems (GWOP) held in Wislikofen, 2006.

2.1 Introduction

In extremal combinatorics the notion of graph decompositions is well-known. Originally, a question by Turán [78] initiated a a search for extremal values of graph parameters for certain classes of graphs. Turán solved the following problem: What is the maximum number of edges in a graph which does not contain a copy of the complete graph K_n as a subgraph? This may be generalized in several ways and we refer to Jukna [42] for further reading on the topic.

Before we present our findings let us mention a related result and an open problem regarding the decompositions of abstract graphs, i.e., omitting our common assumption of a geometric embedding. Any graph on n vertices has a decomposition into at most $\lfloor n/2 \rfloor$ paths and cycles due to Lovász [54]. It is not known whether similar bounds hold when restricting to paths or to cycles only. Alspach [14] proposes the following question: For n odd and $c_1, \ldots, c_k$ natural numbers between 3 and n which sum to $\binom{n}{2}$, does there exist a decomposition of K_n into cycles of length $c_1, \ldots, c_k$?

Here, we consider the complete graph K_n embedded on Γ_n a set of n points in convex position. Our goal is to partition $E(K_n)$ into the smallest possible number of subsets such that each of them induces a crossing-free partition on Γ_n. We show that for n sufficiently large, at least $n - 4$ crossing-free partitions are necessary for that purpose. On the other hand, with partitions that correspond to maximal crossing-free matchings on Γ_n a decomposition of K_n using exactly n partitions is proposed.

This geometric setting allows the following illustration. Suppose n people go to a pub, sit at a round table and order a beer each. Now, traditionally when toasting everyone's glass has to touch every other's exactly once without any crossing occurring, i.e., no-one is allowed to reach under someone else's glass. Surely, we want to minimize the time necessary to complete this procedure. Besides the classical one-on-one "Cheers!" we allow for several people to touch their respective two neighbors' glasses at the same time but forbid any further connection within such a group. Interpreting beer glasses as vertices and the clinking of two of them as an edge then at each round we are considering a crossing-free partition of Γ_n.

2.2 The upper bound

Let Γ_n denote a set of n points in convex position, i.e., the set of vertices of a convex n-gon. Recall that a partition of Γ_n is said to be crossing-free if the convex hulls of the individual parts do not intersect. We will again identify a given crossing-free partition with the straight-line embedded graph on Γ_n containing the boundary edges of the convex hulls of the partition classes. By definition this graph, clearly, is crossing-free.

Theorem 2.1. *For n points in convex position there is a decomposition of K_n into n crossing-free partitions.*

Proof. Let $\{p_1, p_2, \ldots, p_n\}$ be the points of Γ_n given in clockwise order as they appear on the convex hull. The crossing-free partitions used in these decompositions will be large matchings of the underlying point set. To be more precise, in any partition there will be at most two isolated vertices and the remaining vertices will be of degree 1.

Let us first consider the case where n is odd. Then we define an initial matching M_1 of the decomposition by connecting p_i with p_j if and only if $i+j = n+2$; See the left drawing in Figure 2.1. Note that this is a matching where only p_1 is an isolated vertex. Moreover, it is a crossing-free partition: Consider two distinct edges, $p_{i_1}p_{j_1}$ and $p_{i_2}p_{j_2}$, with $i_1 + j_1 = i_2 + j_2$ and $i_1 < i_2$. This implies $j_1 > j_2$, and since the vertices have clockwise order the edges do not cross.

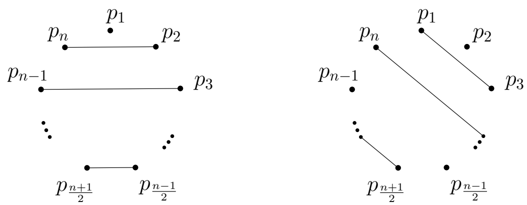

Figure 2.1: Crossing-free matchings M_1 and M_2 on Γ_n, for odd n

For the other crossing-free matchings from the decomposition let $\ell \in \mathbb{N}$, $1 \le \ell \le n-1$, and define

$$M_{\ell+1} := \{p_{i+\ell}p_{j+\ell} \,|\, p_ip_j \in M_1\},$$

where we consider indices modulo n: See the right picture in Figure 2.1 depicting the matching M_2. Similar to the argument given above every such matching is crossing-free. It remains to show that we indeed obtain a decomposition of K_n. Note that all edges used in the construction are distinct: Consider two edges, $p_{i+\ell}p_{j+\ell}$ and $p_{i'+k}p_{j'+k}$, for some ℓ, k, with $i+j = n+2 = i'+j'$. If the edges are the same the sum of their endpoints' indices agrees modulo n, i.e., $i+j+2\ell \equiv i'+j'+2k$ modulo n. This implies $\ell = k$ since n is odd, but all edges in $M_{\ell+1}$ are distinct.

In every matching there is exactly one isolated vertex, hence we constructed n matchings each of size $\frac{n-1}{2}$. Since all these edges are distinct we covered all $n\frac{n-1}{2} = \binom{n}{2}$ edges of K_n in the decomposition.

Now, for n even we consider two kinds of matchings that we will again rotate in order to cover all edges. Let M_1^1 be the matching where we connect vertices p_i and p_j if and only if $i+j=n+1$, and M_1^2 is the matching with edges between p_i and p_j if $i+j=n+2$, see Figure 2.2.

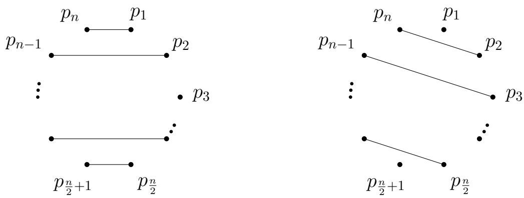

Figure 2.2: Crossing-free matchings M_1^1 and M_1^2 on Γ_n, for even n

By a similar argument as in the case above where n was odd, the two matchings M_1^1 and M_1^2 are crossing-free. Furthermore, since here $n+1$ is odd the indices of an edge p_ip_j in M_1^1 have distinct parity, whereas for an edge in M_1^2 the indices have the same parity. Therefore, all edges of the matchings

$$\begin{aligned} M_{\ell+1}^1 &:= \{p_{i+\ell}p_{j+\ell} \,|\, p_ip_j \in M_1^1\} \\ M_{\ell+1}^2 &:= \{p_{i+\ell}p_{j+\ell} \,|\, p_ip_j \in M_1^2\}, \end{aligned}$$

where $\ell \in \mathbb{N}$ such that $1 \leq \ell \leq \frac{n}{2} - 1$, are distinct. Since M_1^1 has $\frac{n}{2}$ and M_1^2 has $\frac{n-2}{2}$ edges we covered

$$\frac{n}{2} \cdot \frac{n}{2} + \frac{n}{2} \cdot \frac{n-2}{2} = \frac{n}{2} \cdot (n-1) = \binom{n}{2}$$

edges in the decomposition. These are all edges of K_n and we are done. □

2.3 The lower bound

Let us point out that any crossing-free partition has at most n edges, hence in order to cover all $\binom{n}{2}$ edges of the complete graph K_n at least $\frac{n-1}{2}$ partitions are needed. In the following we will show that the true answer about the smallest decomposition of K_n into crossing-free partitions, however, is roughly the upper bound from Theorem 2.1 in the previous section.

Assume once again that the points $\{p_1, p_2, \ldots, p_n\}$ of Γ_n are given in clockwise order as they appear on the convex hull. Then an edge $\{p_i, p_j\}$ of K_n is called *short* if and only if $|i - j| \leq 3$ modulo n, otherwise $\{p_i, p_j\}$ is a *long* edge. In other words, long edges are those for which there are at least three points between p_i and p_j on both arcs of the convex hull of Γ_n. Moreover, given a crossing-free partition Q we denote by $e_s(Q)$ the number of short edges in Q, and similarly we write $e_l(Q)$ for the number of long edges in Q. We identify a crossing-free partition of Γ_n with its canonically induced plane graph.

Note that by definition for $n \leq 7$ all edges of K_n are short, and any crossing-free partition of Γ_8 can have at most one long edge. The following is the key lemma relating the number of long and short edges in a partition.

Lemma 2.2. *Let Q be a crossing-free partition on Γ_n, where $n \geq 7$. Then*

$$e_l(Q) \leq \frac{n + e_s(Q)}{2} - 3.$$

Proof. We give a proof by induction on n. As already mentioned, if $n = 7$ there are no long edges, and if $n = 8$ any crossing-free partition has at most one long edge and the statement holds. For $n = 9$ observe that there are at most two long edges, and if there are exactly two then the partition also has at least one short edge. These three cases will serve as the base of our inductive proof.

Therefore, in the following suppose $n \geq 10$ and without loss of generality we may further assume that Q has no connected component that consists of short edges only. Indeed, the edges of such a component would just increase the right side of the estimate, weakening the statement we want to prove. Now, we can distinguish two cases. Either there is a vertex in Q incident to a long and a short edge, or there are no short edges at all in Q.
Case 1: There are vertices p, q, and r such that pq is long and pr is short.

The line through pq subdivides Γ_n into two parts, and we refer to the one containing r as the part *below* pq. It is crucial to observe that deleting the two vertices p and q does not change the property of being long or short for an edge above pq, since pq is long.

Let us first assume that q is also incident to a short edge and call its adjacent vertex s. Now, delete p and q and their adjacent edges in Q and, if possible, connect the vertices r and s by a new edge. Otherwise, if $r = s$ or the edge rs was already present in Q we do not change anything. This way we obtain a new crossing-free partition Q', with the following properties

$$n(Q') = n(Q) - 2 \qquad e_s(Q') \leq e_s(Q) - 1 \qquad e_l(Q') = e_l(Q) - 1.$$

Therefore, by induction

$$\begin{aligned} e_l(Q) &= e_l(Q')+1 \\ &\leq \frac{n(Q')+e_s(Q')}{2}-3+1 \\ &\leq \frac{n(Q)-2+e_s(Q)-1}{2}+1-3 \\ &\leq \frac{n(Q)+e_s(Q)}{2}-3. \end{aligned}$$

Conversely, we now assume that q is incident to another long edge connecting q with s. As before, deleting p and connecting its neighbors q and r does not change anything above pq, and the new edge qr is long since qp and qs were long. We also note that qs remains a long edge in the new partition Q'. Therefore,

$$n(Q')=n(Q)-1 \qquad e_s(Q')=e_s(Q)-1 \qquad e_l(Q')\geq e_l(Q)-1,$$

where equality holds in the last estimate if and only if $r=s$. By induction

$$\begin{aligned} e_l(Q) &\leq e_l(Q')+1 \\ &\leq \frac{n(Q')+e_s(Q')}{2}-3+1 \\ &= \frac{n(Q)-1+e_s(Q)-1}{2}+1-3 \\ &= \frac{n(Q)+e_s(Q)}{2}-3. \end{aligned}$$

Case 2: There are no short edges in Q at all.

Consider a long edge pq with only isolated vertices on one side. Since Q does not contain a short edge there are at least three such isolated vertices. Now, we turn to the other side and let r and s be the points following p in cyclic order, and similarly t and u the points following q. Observe that since pr and qt are not present in Q, as they would constitute short edges, neither pt nor qr are edges in Q. We distinguish two cases depending on the existence of an edge between r and t.

First, assume that rt is not an edge in Q. Then deleting three isolated vertices below pq can only make this edge pq a short edge if it does at all, and the number of long edges decreases by at most 1. Indeed, note that by assumption neither pu nor qs is an edge in Q. We obtain a new crossing-free partition Q' with

$$n(Q')=n(Q)-3 \qquad e_s(Q')\leq e_s(Q)+1 \qquad e_l(Q')\geq e_l(Q)-1.$$

Then by induction

$$\begin{aligned} e_l(Q) &\leq e_l(Q') + 1 \\ &\leq \frac{n(Q') + e_s(Q')}{2} - 3 + 1 \\ &\leq \frac{n(Q) - 3 + e_s(Q) + 1}{2} + 1 - 3 \\ &= \frac{n(Q) + e_s(Q)}{2} - 3. \end{aligned}$$

Finally, suppose rt is an edge in Q. Then p and q are of degree 1 in Q and their deletion yields a new partition Q' with

$$n(Q') = n(Q) - 2 \qquad e_s(Q') = e_s(Q) \qquad e_l(Q') = e_l(Q).$$

Inductively,

$$\begin{aligned} e_l(Q) &= e_l(Q') \\ &\leq \frac{n(Q') + e_s(Q')}{2} - 3 \\ &= \frac{n(Q) - 2 + e_s(Q)}{2} - 3 \\ &= \frac{n(Q) + e_s(Q)}{2} - 3, \end{aligned}$$

which finishes the proof. □

With the previous lemma we can now prove the main result of this chapter. Fix a decomposition $\mathcal{D}$ of K_n into crossing-free partitions, and denote by $d_i := d_i(\mathcal{D})$ the number of partitions in $\mathcal{D}$ with i short edges. The number of partitions in the decomposition $\mathcal{D}$ is $\sum_{i=0}^{\infty} d_i$ for which we want to derive a lower bound. Since any such partition has at most n (short) edges, we may truncate the sum in order to estimate $|\mathcal{D}| = \sum_{i=0}^{n} d_i$.

Theorem 2.3. *Let $n \geq 7$, then for any decomposition $\mathcal{D}$ of K_n into crossing-free partitions on Γ_n*

$$|\mathcal{D}| \geq n - 4 - \frac{24}{n-6}.$$

Proof. Since we assume $n > 6$ the total number of short edges in K_n is $3n$, or equivalently,

$$\sum_{i=0}^{n} i \cdot d_i = 3n. \tag{2.1}$$

There are $\binom{n}{2}$ edges in K_n which need to be covered by some partition. Observe that, for $0 \leq i \leq n$, the number of edges covered by all partitions with i short edges is at most d_i times the maximum number of edges in a partition with i short edges. Therefore,

$$\begin{aligned}
\binom{n}{2} &\leq \sum_{i=0}^{n} d_i \cdot \max_{Q:e_s(Q)=i} \{e_s(Q) + e_l(Q)\} \\
&\leq \sum_{i=0}^{n} d_i \cdot \left(i + \frac{n+i}{2} - 3\right) \\
&= \frac{3}{2}\sum_{i=0}^{n} i \cdot d_i + \left(\frac{n}{2} - 3\right)\sum_{i=0}^{n} d_i,
\end{aligned}$$

using Lemma 2.2. With the previous observation (2.1) we find

$$\sum_{i=0}^{n} d_i \geq \frac{2}{n-6} \cdot \frac{n(n-1) - 9n}{2} = \frac{n^2 - 10n}{n-6},$$

which proves the claim. □

In particular for $n > 30$ we find that the number of crossing-free partitions in a decomposition of K_n is at least $n - 4$.

On the other hand for small values of n we may compute the exact size of the smallest decompositions by brute force. In the following table we summarize the values we know for the smallest decomposition $\mathcal{D}_n$ of K_n, where $n \leq 8$. For K_9 there is a decomposition into eight crossing-free partitions, however, we do not know whether seven suffice.

n	1	2	3	4	5	6	7	8
$\lvert\mathcal{D}_n\rvert$	1	1	1	3	3	5	6	7

It is trivial to see that the smallest decomposition for $n \leq 3$ only needs a single crossing-free partition to cover all edges of K_n. In the case of $n = 4$ and $n = 5$ points the smallest decomposition of K_n consists of three crossing-free partitions.

As shown in Figure 2.3 and Figure 2.4 three partitions suffice for such decompositions. For proving the corresponding lower bounds suppose that two partitions would be enough. Since K_4 has six edges the partition classes would either both contain three edges, or one partition had two the other four edges. Both cases lead to immediate contradictions.

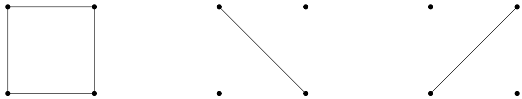

Figure 2.3: Decomposition of K_4 into three crossing-free partitions

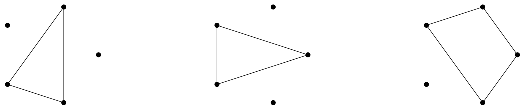

Figure 2.4: Decomposition of K_5 into three crossing-free partitions

Similarly for K_5 containing ten edges the two partition classes both need to have five edges. But there is only one crossing-free partition on Γ_5 with five edges.

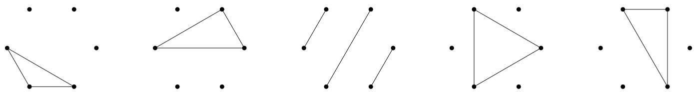

Figure 2.5: Decomposition of K_6 into five crossing-free partitions

To prove the correct lower bound for a decomposition of K_6 a more detailed case analysis is needed. However, since the upcoming examples would also require an even more elaborate proof we refrain from presenting tedious case analyses and refer to the brute force computation of the lower bounds for the smallest decompositions of K_6, K_7, and K_8.

Lastly, for $n = 9$ we do not know the size of the smallest decomposition, however, we can improve over the result from Theorem 2.1 as seen in Figure 2.8.

From the instances drawn above it is unfortunately not clear how to generalize the constructions to yield decompositions of K_n into at most $n - 1$ crossing-free partitions for every n. Already for $n = 10$ it remains open whether nine crossing-free partitions suffice.

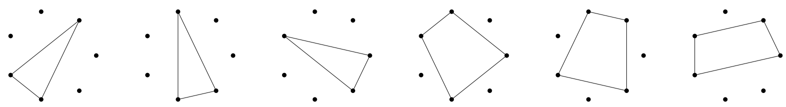

Figure 2.6: Decomposition of K_7 into six crossing-free partitions

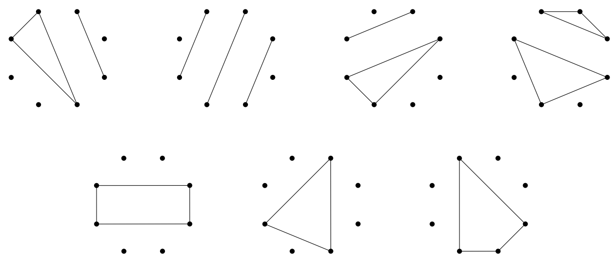

Figure 2.7: Decomposition of K_8 into seven crossing-free partitions

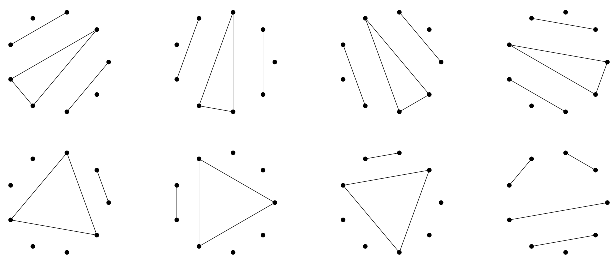

Figure 2.8: Decomposition of K_9 into eight crossing-free partitions

Part II

Transformation Graphs

The truth always lies somewhere else.

Blackboard, Emo Welzl
upper right corner

3

Compatible Spanning Trees

For a planar point set we consider the graph whose vertices are the crossing-free straight-line spanning trees of the given point set, and two such spanning trees are adjacent if their union is crossing-free. An upper bound on the diameter of this graph implies an upper bound on the diameter of the flip graph of pseudo-triangulations of the underlying point set [2].

We prove a lower bound of $\Omega(\log n / \log \log n)$ for the diameter of the transformation graph of spanning trees on a set of n points in the plane. This nearly matches the known upper bound of $O(\log n)$, see [3, 2]. If we measure the diameter in terms of the number k of convex layers of the point set, our lower bound construction is tight, i.e., the diameter is in $\Omega(\log k)$ which matches the known upper bound of $O(\log k)$. So far only constant lower bounds were known.

The results presented in this chapter are joint work with Kevin Buchin, Uli Wagner and Takeaki Uno [22].

3.1 Preliminaries

Given a set P of n points in the plane let $\mathcal{S}t(P)$ denote the set of all crossing-free straight-line spanning trees of P. A straight-line embedded graph is *crossing-free* if no pair of its edges shares any point other than common endpoints. We call two crossing-free spanning trees T_1 and T_2 of P *compatible* if their union, i.e., the graph on P with edge set $E(T_1) \cup E(T_2)$, is crossing-free.

A *tree graph* is a directed graph that has $\mathcal{S}t(P)$ as its vertex set and two vertices (trees) T_1, T_2 are connected by an arc from T_1 to T_2 if the tree T_2 may be obtained from T_1 by some predefined transformation rule. Avis and Fukuda [15] consider the tree graph where two trees are adjacent if their symmetric difference is a path of length 2 which starts at the left-most point of P. Let us note that in this setting the tree graph may be understood as undirected graph. They show that the tree graph is connected and has diameter at most $2n - 4$.

As far as the order of $\mathcal{S}t(P)$ is concerned García et al. [33] prove that the number of crossing-free spanning trees is minimized for a point set Γ_n in convex position, that is when all points lie on the boundary of the convex hull. We will discuss this particular result and related work in more detail in Chapter 5. For the special case of Γ_n, Hernando et al. [37] consider the tree graph where two trees on Γ_n are adjacent if the symmetric difference of their edge sets is of size 2. It is shown that this tree graph is Hamiltonian and has maximum connectivity, which means that its connectivity is equal to the minimum vertex degree. They also give a lower bound of $3n/2 - 5$ for its diameter.

Aichholzer et al. [3] consider the tree graph where the predefined rule is defined by mapping a given tree T in $\mathcal{S}t(P)$ to the tree of minimum Euclidean length which is compatible to T. They show that this tree graph is a rooted tree with the Euclidean minimum spanning tree of P being the root. Furthermore, any tree has distance at most $O(\log n)$ from the root. Another transformation rule which they mention is an operation called *edge slide*, where two trees T and T' are adjacent if there is an edge of T such that keeping one of its endpoints fixed and sliding the other endpoint along a respective adjacent edge yields T'. They show that the tree graph corresponding to this transformation is connected, implying that any two crossing-free spanning trees can be transformed into each other by means of local and constant-size changes only. Recently, Aichholzer and Reinhardt [11] gave an upper bound of $O(n^2)$ for the edge slide distance between any two crossing-free trees.

In this chapter we are interested in the tree graph $\mathcal{T}_{\mathrm{st}}(P)$ with $\mathcal{S}t(P)$ as vertex set and edges between compatible spanning trees. Note that by symmetry of the definition of compatible trees the tree graph can be considered as undirected. Aichholzer et al. [2] refine the upper bound of $O(\log n)$ on the diameter of $\mathcal{T}_{\mathrm{st}}(P)$, given in [3], to a bound of $O(\log k)$, where k denotes the number of convex layers of P. The *convex layers* of a point set P are defined inductively: The first convex layer U_1 consists of the extreme points of the convex hull of P, and for $i > 1$ the i-th convex layer U_i is defined as the set of extreme points on the boundary of the convex hull of $P \setminus \bigcup_{j<i} U_j$. The number k of convex layers is the minimum i such that $U_{i+1} = \emptyset$.

Consider a plane polygon with vertices in P of which exactly three are convex, then the bounded face is called *pseudo-triangle on* P. A *pseudo-triangulation* of a given point set P is a plane graph where every face is a pseudo-triangle. The flip graph of pseudo-triangulations of P is defined as the graph whose vertices are the pseudo-triangulations of P with edges between pseudo-triangulations that differ in exactly one edge, either by replacement or by removal. Aichholzer et al. [2] prove that an upper bound of d on the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ yields an upper bound of $O(nd)$ on the diameter of the flip graph of pseudo-triangulations of P.

Another related problem, transforming compatible perfect matchings, was very recently treated by Aichholzer et al. [6]. The notion of compatible perfect matchings is defined analogously to that of spanning trees. In their work [6] it is shown that a sequence of length $O(\log n)$ suffices to transform a given perfect matching into any other perfect matching on a fixed set of $n = 2m$ points. This improves the previously best known linear upper bound of $n - 2$ by Houle et al. [38], who were the first to show that the corresponding transformation graph is connected. In Chapter 4 we will prove a corresponding lower bound of $\Omega(\log n/ \log\log n)$.

The transformation graph for the special case of convex position, and with two perfect matchings being adjacent if they differ in exactly two edges, is known to be bipartite and of diameter $n-2$, which is a result due to Hernando et al. [36]. Moreover, they show that this transformation graph is Hamiltonian if m is even, and the graph does not contain a Hamiltonian path for odd m.

In [2] it is conjectured that the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ is sub-logarithmic, with then no example known where the diameter is not constant. We give a sub-logarithmic but considerably tight lower bound by complementing the $O(\log n)$ upper bound with a lower bound of $\Omega(\log n/ \log\log n)$. This is achieved constructively by providing point sets of increasing size, and on each point set we specify two spanning trees achieving this bound. We also present an example for which the bound in the number of convex layers is tight, i.e., the distance between the two trees is $\Omega(\log k)$, where k is the number of convex layers of the underlying point set.

3.2 The lower bound

In this section we construct point sets in the plane and consider pairs of spanning trees which need a large number of transformation steps to transform one tree into the other.

We will first develop a general scheme to construct such trees. Based on this we present two recursive constructions using the scheme in different ways. The first construction yields a lower bound of $\Omega(\sqrt{\log n})$ for the number of transformations, where n is the size of the underlying point set. However, in terms of the number k of convex layers the lower bound is tight, i.e., the diameter is $\Omega(\log k)$. The second construction gives a lower bound of $\Omega(\log n / \log\log n)$. For the sake of simplicity of the description we initially use point sets with more than two points on a line, i.e., the points are not in general position. However, they can easily be changed to do so by applying a small perturbation without losing any of the relevant properties of the construction. We comment on these perturbations later on.

The key idea of both constructions is to place the top-most vertex of the point set very far away from the others. We will consider a first tree with many nearly horizontal edges and a second tree with only nearly vertical edges crossing many of the horizontal edges in the first tree. Furthermore, there are dependencies between the horizontal edges. During the transformation a vertex incident to a horizontal edge may connect to the top-most vertex by a vertical edge only if certain horizontal edges are no longer in the current tree.

We illustrate this by the example in Figure 3.1 with $P = \{a, b, \ldots, h\}$ being the underlying point set. The first tree T_1 in Figure 3.1(a) consists of mostly nearly horizontal edges, whereas the second tree T_2 in Figure 3.1(b) has only nearly vertical edges. The points b, c, d, e, and f define two vertical strips containing all points of P. In each such strip there is a point at the bottom (g and h, respectively) which needs to connect to the top-most point a through the corresponding strip, see Figure 3.1(c). At the beginning of any transformation the edges bc and ef block both strips completely, i.e., both the bottom-most points g and h cannot connect to a in any neighbor

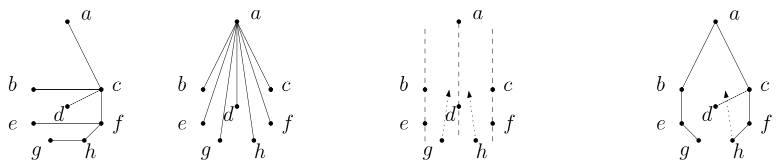

(a) T_1 (b) T_2 (c) Two vertical strips defined by b, c, d, e, f (d) One vertical strip is blocked

Figure 3.1: The trees T_1 and T_2 have distance 3 in $\mathcal{T}_{\text{st}}(P)$

of T_1 in $\mathcal{T}_{\text{st}}(P)$. Whatever the first transformation is, thereafter the point d will be adjacent to at least one of b, c, e, or f; For instance, we obtain the tree in Figure 3.1(d). Thus, after one transformation the edge ag or ah still crosses an edge of the current tree and cannot be present after the next transformation. Hence, three transformations are necessary, and also suffice, to transform T_1 to T_2, and the diameter of $\mathcal{T}_{\text{st}}(P)$ is at least 3.

Blocking vertical strips

Before turning to the general constructions of the point sets, we further develop the concept of blocking vertical strips. A *vertical strip* R is a subset of $\mathbb{R}^2$ such that there exist $\alpha, \beta \in \mathbb{R}$, $\alpha \leq \beta$, with

$$R = \left\{(x, y) \in \mathbb{R}^2 \middle| \alpha \leq x \leq \beta\right\} = [\alpha, \beta] \times \mathbb{R};$$

The *width* of the vertical strip R is $\beta - \alpha$. An edge *blocks a vertical strip* if the endpoints of the edge lie on different sides or possibly on the boundary of the strip. For the continuation of our example from Figure 3.1(a) we assume in the following a coordinate system with the point b at coordinates $(0, 3)$, c at $(1, 3)$, d at $(1/2, 2)$, e at $(0, 1)$, and f at $(1, 1)$. Then the edges bc and ef of the tree T_1 both block the vertical strip $[0, 1] \times \mathbb{R}$ and the edge dc blocks the vertical strip $[1/2, 1] \times \mathbb{R}$.

A point set S together with a set E of straight-line edges on S *blocks a vertical strip of width* $w > 0$ *after* k *steps*, if for any point set P containing S but no further point inside the convex hull of S the following holds: If a spanning tree $T \in \mathcal{S}t(P)$ contains the edges E then in any spanning tree from the k-neighborhood of T in $\mathcal{T}_{\text{st}}(P)$ some vertical strip of width at least w is blocked, not necessarily by an edge in E. A corresponding point set P is said to have the *blocking property*. For instance, in the tree T_1 in Figure 3.1(a) the points b, c, d, e, and f together with the edges bc and

ef block a vertical strip of width $1/2$ after one step, since either the strip $[0, 1/2] \times \mathbb{R}$ is blocked by bd or ed, or the strip $[1/2, 1] \times \mathbb{R}$ is blocked by dc or df.

Note that this concept now implies the following. Assume that we have a point set P with the top-most point $p_0 \in P$ placed very far away from the rest, and a subset $S \subset P$ with edges E on S blocks some vertical strip R after k steps. Let $T_1 \in \mathcal{S}t(P)$ be a tree containing the edges E and let $T_2 \in \mathcal{S}t(P)$ be the tree where p_0 connects to every other point in P by a nearly vertical edge. If there is a point in $P \cap R$ lying strictly below the edge blocking R after k steps then T_2 cannot be in the $(k+1)$–neighborhood of T_1 in $\mathcal{T}_{\mathrm{st}}(P)$. Thus, the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ is at least $k+2$.

The point sets we are about to construct all reside in the strip $[0, 1] \times \mathbb{R}$, and therein we consider specific vertical strips that might be blocked. We call a point set S together with a set E of edges an *ℓ-of-m-blocker after k steps* if S blocks at least ℓ of the m vertical strips $[(i-1)/m, i/m] \times \mathbb{R}$ after k steps, for $1 \leq i \leq m$. We point out that for distinct trees containing E we do not require the same strips to be blocked in their respective k–neighborhood. We call ℓ/m the *density* of the blocker.

In the example of T_1 in Figure 3.1(a) the points b, c, d, e, f together with the edges bc and ef are a 1-of-2-blocker after one step. We will in the following refer to the point set of this particular blocker as A. By stacking enough copies of a given blocker and spreading further points in-between we can construct new blockers with an increased number of steps.

We will present two ways of using this principle. First we build blockers of small density at most $1/2$, i.e., for $d \geq 1$ we obtain a 1-of-2^d-blocker after d steps. In the second construction we build blockers keeping the density as large as possible.

The advantage of the first construction is that the blocker only requires an exponential number of vertical strips, and we obtain a tight bound of $\Omega(\log k)$ for the diameter of $\mathcal{T}_{\mathrm{st}}(P)$, where k is the number of convex layers of P. However, the number of previously constructed blockers needed in the recursive construction will also grow exponentially. This yields a blocker after d steps with $O\big(2^{d^2}\big)$ rows, and for the corresponding point set P containing this blocker we have $d = \Omega(\sqrt{\log n})$, where $n = |P|$.

In the second construction we will only use a number linear in d of previously constructed blockers. For this construction to work, we spread in more points horizontally in each step. This will result in the number of rows being of order $O(d^d)$, and we arrive at the same super-exponential bound for the number of vertical strips. The construction gives a point set P such that the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ is $\Omega(\log n / \log\log n)$.

3.2.1 Construction 1

We begin by extending the previous example of the 1-of-2-blocker after one transformation step.

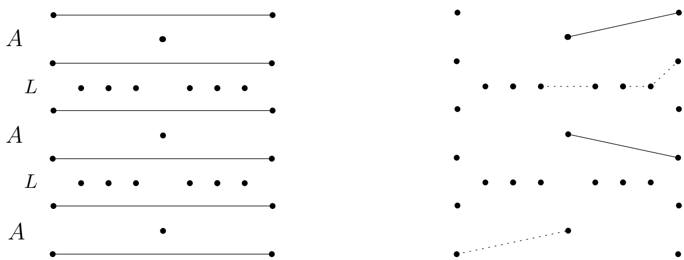

(a) Stacking three 1-of-2-blockers, with additional points between the blockers

(b) Pair of edges (solid lines) blocking the same strip after one step

Figure 3.2: Construction of a 3-of-8-blocker after two steps

Consider the construction given in Figure 3.2(a). It contains three copies of the 1-of-2-blocker after one step, i.e., the point set A together with the corresponding horizontal edges, and between two subsequent blockers there is a copy of an additional point set L. Note that L subdivides each of the two strips of width 1/2 into four smaller strips resulting in a total of eight vertical strips. Each copy of A blocks one vertical strip of width 1/2 after one step. Since there are three copies of A by the pigeon-hole principle one strip is blocked twice. In the example of Figure 3.2(b) this is the right vertical strip. No matter how the points of L between these blocking edges are connected to the rest of the tree at least three of the four corresponding vertical strips of width 1/8 are blocked. This can only change after the edges blocking the strip of width 1/2 are removed. For that to happen at least one more step is required, thus the construction is a 3-of-8-blocker after two steps. As described before, choosing a point set P and a tree on P containing the blocker from Figure 3.2(a) results in the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ being at least 4.

Lemma 3.1. *Let S be an ℓ-of-m-blocker after k transformations with density $\ell/m > 1/u$, for some $u \in \mathbb{N}$. By stacking u copies of S on top of each other and placing additional points between each pair of subsequent copies that equidistantly subdivide each of the m vertical strips into m' smaller strips we obtain an $(m'-1)$-of-$(m \cdot m')$-blocker after $k+1$ steps.*

Proof. After k steps the u copies of S block within the m vertical strips $\ell{\cdot}u > m$ times, thus at least one of the m strips is blocked twice. The points in this vertical strip blocked from above and below subdivide this strip into m' smaller strips, hence in order to connect these points to the rest at least $m' - 1$ of the small strips are blocked. This changes at the earliest after $k+1$ steps, thus the construction is an $(m'-1)$-of-$(m \cdot m')$-blocker after $k+1$ steps. □

As a first implication, Lemma 3.1 readily shows that the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ cannot be bounded by a constant for arbitrary point sets P in the plane. To be more specific, given $d \in \mathbb{N}$, we will in the following construct a point set P together with two trees $T_1, T_2 \in \mathcal{S}t(P)$ such that at least d steps are necessary to transform one of the trees into the other, and the size of P is in $O\big(2^{d^2}\big)$, i.e., $d = \Omega(\sqrt{\log n})$, where $n = |P|$.

All points of P lie in the infinite strip $[0,1] \times \mathbb{R}$. A special point $p_0 \in P$ has a larger y-coordinate than all other points, and will be chosen such that the slope of any line through p_0 and an other point in P is larger than the slopes of all non-vertical lines through two points from $P \setminus \{p_0\}$.

Let us define point sets $L_0 := \{(0,0),(1,0)\}$ and L_k, for $k \in \mathbb{N}$, as

$$L_k := \left\{ \Big(\frac{2i-1}{2^k}, 0\Big) \,\middle|\, 1 \le i \le 2^{k-1} \right\}.$$

Thus, $\bigcup_{0\le\ell\le k} L_\ell$ contains 2^k+1 points and subdivides the line segment from $(0,0)$ to $(1,0)$ into 2^k equal parts. The set L_{k+1} places one point in the center of each of these parts.

We also define point sets A_k, $k \in \mathbb{N}$, inductively. Let

$$A_1 := L_0 \;\cup\; L_1 \oplus_y 1 \;\cup\; L_0 \oplus_y 2,$$

where $S \oplus_y i := \{(x, y+i) | (x,y) \in S\}$ is a vertical shift of the point set $S \subseteq \mathbb{R}^2$ by $i \in \mathbb{N}$. Note that A_1, see the left drawing in Figure 3.3, corresponds to the point set A which we already encountered in previous examples.

For $k \ge 1$, let A_{k+1} be defined by stacking 2^k+1 copies of A_k with a copy of L_{k+1} between each pair of subsequent copies of A_k. Formally,

$$A_{k+1} := \bigcup_{i=0}^{2^k} A_k \oplus_y (i \cdot h_k) \;\cup\; \bigcup_{i=1}^{2^k} L_{k+1} \oplus_y (i \cdot h_k - 1),$$

where $h_k := 4 \prod_{i=1}^{k-1} (2^i + 1)$.

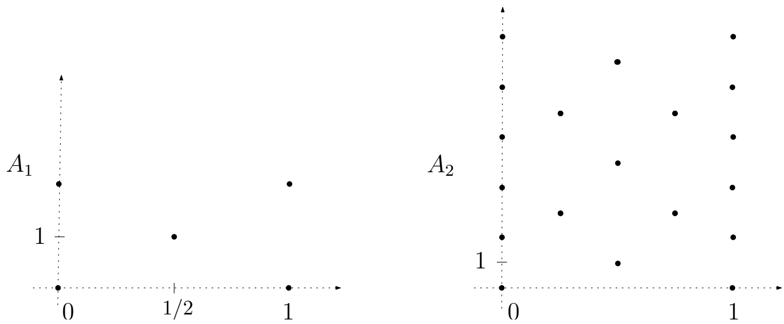

Figure 3.3: Point sets A_1 and A_2 used in the recursive construction

It follows directly from Lemma 3.1 that the point set A_k together with edges between every pair of points with coordinates $(0, y)$ and $(1, y)$, for some $y \in \mathbb{N}$, is a 1-of-2^k-blocker after k steps.

Given $d \in \mathbb{N}$, define $P := L_{d+1} \cup A_d \oplus_y 1 \cup \{p_0\}$ with p_0 chosen as described above. Let T_1 be a tree on P that contains all (exactly) horizontal edges blocking the complete vertical strip $[0, 1] \times \mathbb{R}$ and arbitrarily add further edges such that T_1 becomes a crossing-free straight-line spanning tree. Define T_2 to be the star connecting p_0 to every other point by an edge. We already know that A_d together with the corresponding horizontal edges is a 1-of-2^d-blocker after d steps. Thus, when transforming T_1 into T_2 there will be one of the points in L_{d+1} blocked away from p_0 after d steps. Therefore, at least $d + 2$ transformations are needed.

The cardinality s_d of the point set A_d is given recursively by $s_1 = 5$ and the identity

$$s_{k+1} = (2^k + 1)s_k + 2^k \cdot 2^k.$$

Generously estimating we get $s_{k+1} \le 2^{2k} s_k + 2^{2k} s_k = 2^{2k+1} s_k$ and by induction $s_d \le 5 \cdot 2^{d^2}$ follows. The size of P is $n = 2^d + s_d + 1 = O(2^{d^2})$, hence $d = \Omega(\sqrt{\log n})$.

To conclude we want to express the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ in terms of the number of convex layers. The first layer of P consists of the top-most point p_0, the points of the bottom row, the points in the left-most and the right-most column of the construction. With each additional convex layer two more rows and two further columns are considered until only one row or one column is left. If m_1 is the number of different x-coordinates and m_2 the number of distinct y-coordinates used in the construction then we can bound the number of convex layers from above by

$$1 + \frac{1}{2} \min(m_1, m_2).$$

The number of distinct x-coordinates in P is $2^{d+1} + 2$ and thus d is logarithmic in the number of convex layers. This bound is tight as shown by the result of Aichholzer et al. [2] which we provide in Section 3.3. Note that the number of distinct y-coordinates is $h_d + 1$ which is of order $O\big(2^{d^2}\big)$.

At this point we want to mention that perturbing P slightly does not destroy the blocking property. In order to see this we observe that the shortest blocking edge after d steps still has length at least 2^{-d}. It is this nearly horizontal edge that prevents some point in L_{d+1} from connecting to p_0 in the next transformation step. Now, choose $0 < \varepsilon \ll 2^{-(d+1)}$ and slightly perturb every point in P by at most ε such that the resulting point set is in general position. The width of any strip we encounter during a transformation decreases by at most 2ε which is negligible compared to its original size. Hence, also the perturbed point set has a blocking edge after d transformation steps. In order for the derived lower bound to hold as well in terms of the number of convex layers we need that points on a convex layer remain there even after the perturbation.

Theorem 3.2. *For any sufficiently large integer k there is a set P of points in general position in the plane which consists of k convex layers such that the diameter of $\mathcal{T}_{\text{st}}(P)$ is $\Omega(\log k)$. This diameter is also $\Omega(\sqrt{\log n})$, where n is the number of points in P.*

3.2.2 Construction 2

The number of distinct x-coordinates of points in A_k from Construction 1 was growing exponentially with d. On the one hand, this implied that the diameter of the corresponding transformation graph is at least logarithmic in the number of convex layers. On the other hand, A_k suffered from an exponential growth in the number of copies of previously constructed sets A_{k-1}, resulting in a doubly-exponential growth in the number of points. The recursive construction we present in the following will only require a linear number of copies of previously considered point sets.

To this end we construct a point set $P := P_d \subset [0, 1] \times \mathbb{R}$ depending on an integer variable $d > 1$, and specify two trees $T_1, T_2 \in \mathcal{S}t(P)$ such that $d = \Omega(\log n / \log\log n)$, where n is the size of P, and the distance of the trees in $\mathcal{T}_{\text{st}}(P)$ is at least d.

Again, we will define point sets $L_k = L_k(d)$ that equidistantly subdivide vertical strips and recursively construct blockers $A_k = A_k(d)$ with similar meaning as in Construction 1. Finally, we also include a special point p_0 in P with a far larger y-coordinate than any other point in P.

However, contrary to the first construction where the density of the blockers dropped by a factor of 1/2 in every step, the density will now only decrease linearly, hence the dependencies of L_k and A_k on d. This is achieved by much denser sets L_k such that in each step the number of vertical strips will grow by a factor of d, instead of doubling.

The key to make this approach work is the following observation. Let $L_0 := \{(0,0),(1,0)\}$ and define

$$A := L_0 \cup \left\{ \left(\frac{i}{d}, 1 \right) \mid 1 \leq i \leq d-1 \right\} \cup L_0 \oplus_y 2.$$

The set A together with the horizontal edges connecting the points in L_0 and $L_0 \oplus_y 2$ constitutes a $(d-1)$-of-d-blocker after one step. In other words, there is only one single strip of width $1/d$ which is not blocked after the first step. Stacking, for instance, three copies of A on top of each other implies that after the first step there cannot be more than one vertical strip of width $1/d$ which is not blocked at least twice. Assume we equidistantly subdivide each of the d strips in A further into d smaller strips by adding additional points, and let us call the resulting point set B. Figure 3.4(a) shows the corresponding construction for $d = 4$ together with the horizontal edges needed for the blocking property.

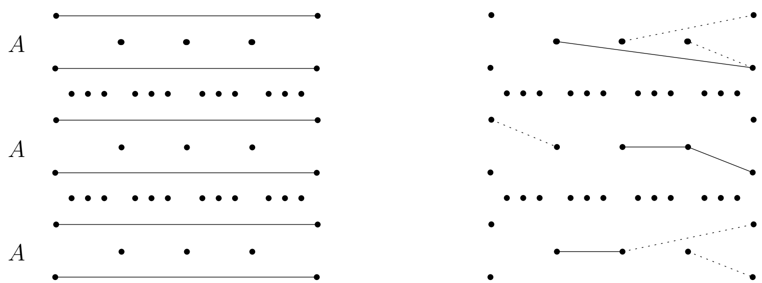

(a) The blocker B, for $d = 4$, consisting of three 3-of-4-blockers A and additional points in-between

(b) Edges blocking the same strips (drawn as solid lines) after one step

Figure 3.4: Construction of a 9-of-16-blocker after two steps

Then each of the $d-1$ vertical strips of width $1/d$ which is blocked twice after the first transformation, together with the additional points in-between, behaves like a (horizontally) scaled and slightly perturbed blocker A for the following transformation step, see Figure 3.4(b). Thus, B together with the corresponding horizontal edges is a $(d-1)^2$-of-d^2-blocker after two steps.

For an even more general setup, let $a_1 \in \mathbb{N}$ be the number of the initial subdivisions of the vertical strip $[0,1] \times \mathbb{R}$ and accordingly define

$$\begin{aligned} L_1 &:= \left\{ \left(\frac{i}{a_1}, 0 \right) \mid 1 \le i \le a_1 - 1 \right\} \\ A_1 &:= L_0 \,\cup\, L_1 \oplus_y 1 \,\cup\, L_0 \oplus_y 2. \end{aligned}$$

Let $b_1 \in \mathbb{N}$ be the minimum number of vertical strips blocked in A_1 after one step, thus $b_1 = a_1 - 1$ since A_1 is an (a_1-1)-of-a_1 blocker after one step. We use A_1 as the base gadget in the upcoming construction.

For $k \ge 1$, we construct A_{k+1} recursively by vertically stacking $c_{k+1} \in \mathbb{N}$ many copies of A_k, and in-between two such consecutive copies we equidistantly subdivide each previous vertical strip into $a_{k+1} \in \mathbb{N}$ smaller strips. Formally, let $\pi_\ell = \prod_{i=1}^{\ell} a_i$ denote the number of strips in A_ℓ, where $\ell \in \mathbb{N}$, then we define

$$L_{k+1} := \left\{ \left(\frac{i}{\pi_k} + \frac{j}{\pi_{k+1}},\ 0 \right) \,\middle|\, \begin{array}{c} 1 \le i \le \pi_k - 1 \\ 1 \le j \le a_{k+1} - 1 \end{array} \right\},$$

i.e., L_{k+1} contains $\pi_k(a_{k+1}-1)$ points which subdivide each vertical strip given by $\bigcup_{0 \le \ell \le k} L_\ell$ into a_{k+1} equidistant parts along the line segment from $(0,0)$ to $(1,0)$. Moreover, for $k \ge 1$, we set

$$A_{k+1} := \bigcup_{i=0}^{c_{k+1}-1} A_k \oplus_y (i \cdot h_k) \ \cup \ \bigcup_{i=1}^{c_{k+1}-1} L_{k+1} \oplus_y (i \cdot h_k - 1),$$

where $h_k := 4 \prod_{i=2}^{k-1} c_i$.

We will denote by $b_k \in \mathbb{N}$ the minimum number of strips of width π_k^{-1} that are blocked in A_k after any k steps of transformation. Notice that a_k, b_k, and c_k might depend on d, and for certain choices of the sequences a_k and c_k the corresponding sets A_k need not be blockers. In order for the construction to work during at least d steps, i.e., so we obtain a set A_d in which there is at least one strip blocked after any d transformations, it is necessary and sufficient to require $b_k > 0$, for all $1 \le k \le d$.

There are sequences meeting this criterion as seen by Construction 1. For the purpose of computing the density of the blockers let us now assume that we are given two such sequences $a_k \geq 2$ and $c_k \geq 2$ yielding sets A_k which together with the corresponding horizontal edges are blockers after k steps.

Lemma 3.3. *The lower bound b_ℓ on the number of strips of width π_ℓ^{-1} that are blocked in A_ℓ after ℓ transformation steps is given by the recursion*

$$\begin{aligned} b_1 &= a_1 - 1 \\ b_{k+1} &= \left\lceil \frac{b_k c_{k+1} - \pi_k}{c_{k+1} - 1} \right\rceil \cdot (a_{k+1} - 1), \quad k \geq 1. \end{aligned}$$

Proof. We prove this statement by induction on the number k of transformation steps. We already discussed the case of our base gadget A_1 where after the first transformation at most one vertical strip of width $1/a_1$ is not blocked, hence $b_1 = a_1 - 1$.

For the inductive step assume that in A_k at least b_k strips of width π_k^{-1} are blocked after k steps and, correspondingly, at most $\pi_k - b_k$ strips are not blocked. When vertically stacking c_{k+1} many copies of A_k in order to obtain A_{k+1} every copy yields at most $\pi_k - b_k$ strips which lack a blocking edge after k steps, hence, in total at most $(\pi_k - b_k) \cdot c_{k+1}$ blocking edges are missing. Now, imagine an adversary distributes these missing blocking edges to the strips of width π_k^{-1} so that there are as many strips as possible which are blocked at most once. Clearly, then for each such strip at least $c_{k+1} - 1$ of the missing blocking edges have to be used. Therefore, after k transformation steps at most $\left\lfloor (\pi_k - b_k) \cdot \frac{c_{k+1}}{c_{k+1}-1} \right\rfloor$ strips of width π_k^{-1} are blocked at most once. The remaining strips, of which there are at least

$$\pi_k - \left\lfloor (\pi_k - b_k) \cdot \frac{c_{k+1}}{c_{k+1} - 1} \right\rfloor = \left\lceil \frac{b_k c_{k+1} - \pi_k}{c_{k+1} - 1} \right\rceil,$$

are blocked twice and thus behave like a horizontally scaled base gadget for the next transformation. Therefore, A_{k+1} is a b_{k+1}-of-π_{k+1} blocker after $k+1$ steps, with b_{k+1} as claimed in the recursion. □

Since $a_k, c_k \geq 2$, for all k, the condition $b_d > 0$ for the target blocker A_d is equivalent to $b_{d-1} c_d - \pi_{d-1} > 0$, which in particular implies $b_{d-1} > 0$. Recursively $b_k > 0$ follows, for the remaining $1 \leq k \leq d-2$, and hence $b_d > 0$ is a necessary and sufficient condition for the construction to yield a blocker after d steps.

For computational purposes we drop the requirement on the b_k being natural numbers and use the following slightly weaker recursion than in the statement of Lemma 3.3

$$\begin{aligned}\tilde{b}_1 &:= a_1 - 1\\ \tilde{b}_{k+1} &:= (\tilde{b}_k c_{k+1} - \pi_k) \cdot \frac{a_{k+1} - 1}{c_{k+1} - 1}.\end{aligned} \tag{3.1}$$

For this new sequence $\tilde{b}_k \leq b_k$ follows readily by induction, and in particular $\tilde{b}_d > 0$ is sufficient for the construction to work.

Now, $\tilde{b}_d > 0$ together with the definition (3.1) of the sequence $\tilde{b}_k$ is equivalent to $\tilde{b}_{d-1} \cdot c_d > \pi_{d-1}$, and hence to

$$\begin{aligned}\frac{\pi_{d-1}}{c_d} &< \tilde{b}_{d-1} = (\tilde{b}_{d-2} c_{d-1} - \pi_{d-2}) \cdot \frac{a_{d-1} - 1}{c_{d-1} - 1},\\ \tilde{b}_{d-2} &> \frac{\pi_{d-1}}{c_d} \cdot \frac{c_{d-1} - 1}{c_{d-1}(a_{d-1} - 1)} + \frac{\pi_{d-2}}{c_{d-1}}.\end{aligned}$$

Recursively plugging in (3.1) for $\tilde{b}_k$ we find

$$\tilde{b}_1 > \sum_{i=0}^{d-2} \frac{\pi_{d-1-i}}{c_{d-i}} \cdot \prod_{j=2}^{d-1-i} \frac{c_j - 1}{c_j(a_j - 1)} = \sum_{i=2}^{d} \frac{\pi_{i-1}}{c_i} \cdot \prod_{j=2}^{i-1} \frac{c_j - 1}{c_j(a_j - 1)}.$$

Since $\tilde{b}_1 = a_1 - 1$ and with the definition of π_ℓ we obtain

$$1 - \frac{1}{a_1} > \sum_{i=2}^{d} \frac{1}{c_i} \cdot \prod_{j=2}^{i-1} \frac{(c_j - 1)a_j}{c_j(a_j - 1)}, \tag{3.2}$$

as a sufficient condition for the sequences a_k and c_k to yield a well-defined construction of a blocker A_d after d transformation steps.

Let us now calculate the size of the point sets A_k, and for this purpose denote $n_k := |A_k|$, for $k \geq 1$, and recall $n_1 = a_1 + 3$.

Lemma 3.4. *For $k \geq 2$, the number of points in A_k is*

$$n_k = n_1 \prod_{i=2}^{k} c_i + \sum_{i=2}^{k} \pi_{i-1}(a_i - 1)(c_i - 1) \prod_{j=i+1}^{k} c_j.$$

Proof. The induction base holds since $n_2 = n_1 c_2 + (c_2 - 1)a_1(a_2 - 1)$. For the inductive step recall that A_{k+1} is constructed by stacking c_{k+1} copies of A_k and subdividing the strips by spreading points in-between the copies. Thus,

$$n_{k+1} = c_{k+1}n_k + (c_{k+1} - 1)\pi_k(a_{k+1} - 1),$$

and substituting the inductive hypothesis for n_k proves the statement. $\square$

The values $a_1 = d + 1$ and $a_k = c_k = d$, for all $2 \leq k \leq d$, satisfy (3.2)

$$1 - \frac{1}{d+1} > \frac{d-1}{d} = 1 - \frac{1}{d},$$

hence we obtain a blocking set A_d after d transformation steps. With $P := L_{d+1} \cup A_d \oplus_y 1 \cup \{p_0\}$, for $a_{d+1} = 2$, and T_1, T_2 defined as in the first construction, the distance of the two trees in $\mathcal{T}_{\mathrm{st}}(P)$ is at least $d + 2$.

By Lemma 3.4 the cardinality n of the point set P is

$$n = 2d^d + \big((d+4)d^{d-1} + d^{d-1}(d-1)^3\big) + 1 = O\big(d^{d+2}\big),$$

and hence $d = \Omega(\log n / \log\log n)$.

Similarly to Construction 1 we may perturb the points by a small ε, with $0 < 2\varepsilon \ll \pi_d^{-1}$, such that P is in general position and neither the blocking property is destroyed nor the convex layers of P are changed. The width of the smallest blocked strip of the perturbed $\tilde{A}_d$ will be at least $\pi_d^{-1} - 2\varepsilon \gg 0$, hence, after d transformations there still is a blocking edge implying that the distance of the corresponding trees in $\mathcal{T}_{\mathrm{st}}(P)$ is $d + 2$.

Theorem 3.5. *For any sufficiently large integer n there exists a set P of n points in general position in the plane for which the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ is $\Omega(\log n / \log\log n)$.*

There are $h_k - 1 = 4\prod_{i=2}^{k-1} c_i - 1 \leq n_k$ distinct y-coordinates in A_k, and also the number $\pi_k + 1$ of x-coordinates is at most n_k. In particular our choice for a_k and c_k implies that the number of convex layers in P is also $O(d^d)$, hence, this construction is not tight in terms of the convex layers.

We have the feeling that the $1/\log\log n$ factor in the lower bound of Theorem 3.5 is more likely to be an artifact of our construction than the truth about the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ which we think should be $\Theta(\log n)$ for a suitable point set P. To construct such a set in the spirit of our approach it would be sufficient to provide sequences a_k, c_k satisfying (3.2) where neither linearly many a_k nor c_k are of linear order in d. The reason why we present a more general framework where we allow for a_k and c_k to take different values is that otherwise (3.2) and the geometric-harmonic means inequality immediately yields

$$\sqrt[d]{n_d} \geq \sqrt[d]{\prod_{i=2}^{d} c_i} \geq \frac{d}{\sum_{i=2}^{d} c_i^{-1}} > d,$$

and in turn $n_d = \Omega(d^d)$, and we obtain the same bound as in Theorem 3.5.

On the other hand, despite considerable effort we could not prove that the general approach will fail to provide a better estimate. Similarly to definition (3.1) one may infer an upper bound on the sequence b_k as given by the following recursion

$$\begin{aligned} b'_1 &:= a_1 - 1 \\ b'_{k+1} &:= \Big((b'_k + 1)c_{k+1} - \pi_k\Big) \cdot \frac{a_{k+1} - 1}{c_{k+1} - 1}. \end{aligned}$$

It is easily checked that indeed $b_k \leq b'_k$ holds, for all k, and hence the necessary condition $b_d > 0$ implies $b'_d > 0$. From this in turn follows

$$1 - \frac{1}{a_1} > \sum_{i=2}^{d} \left(\frac{1}{c_i} - \frac{1}{\pi_{i-1}}\right) \cdot \prod_{j=2}^{i-1} \frac{(c_j - 1)a_j}{c_j(a_j - 1)},$$

as a necessary condition for the construction to work. It remains open whether this already implies an asymptotic behavior of order $\Omega(d^d)$ for n_d.

3.3 The upper bound

Concluding this chapter we will prove the upper bound of $O(\log k)$ on the diameter of the transformation graph of compatible spanning trees on a given point set with k convex layers. The original proof of this statement presented in [2] lacks a small argument which we shall provide in the course of the description, but otherwise our demonstration closely follows the original lines.

Recall the recursive definition of the convex layers for a point set P: The first layer U_1 consists of the extreme points of the convex hull of P and, for $i > 1$, the i-th layer U_i is the set of extreme points of the convex hull of $P \setminus \bigcup_{j<i} U_j$. The number k of convex layers of P is the minimum i such that $U_{i+1} = \emptyset$.

Trivially, a set P of n points has at most n convex layers, hence the upper bound $O(\log n)$ for the diameter of $\mathcal{T}_{\text{st}}(P)$ follows once we proved the bound in terms of the number of convex layers. First, we make the following useful observation.

Lemma 3.6 (Aichholzer et al. [2]). *Consider a triangulation on P, and let $p \in P$ be on some layer U_i, for $i \geq 2$. Then the triangulation contains an edge pq that does not cross U_i such that q lies on a layer U_j, with $j < i$.*

Proof. If such an edge does not exist then all edges incident to p lie within an angle of π. However, the only points in a triangulation with a reflex angle are the ones on the first layer U_1. □

We need some more definitions for the proof of the final result. Consider the edges representing the boundary of the convex hull of a layer U_i which we refer to as *layer edges*. Then a *layer tree* of P is a crossing-free spanning tree on P that, for $i \geq 1$, contains all but one layer edges from each convex layer U_i of P and which connects consecutive layers by single edges that do not cross any layer edge whether in the tree or not, see Figure 3.5 for an example. The last property of layer trees will turn out to be necessary in the proof of the main theorem.

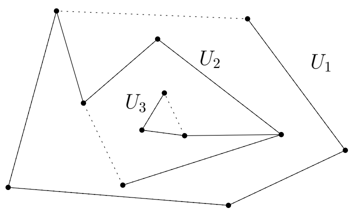

Figure 3.5: A layer tree for a point set with $k = 3$ layers

For any point set P such layer trees exist as a consequence of Lemma 3.6. Indeed, consider a triangulation containing all layer edges of P, then for $i \geq 2$ we can choose an edge connecting U_i and U_{i-1} without crossing U_i. Together with the layer edges chosen as described above the resulting graph is crossing-free and connected, in particular a layer tree of P.

We claim that two layer trees of P have distance at most 2 in $\mathcal{T}_{\text{st}}(P)$. Two such trees clearly do not intersect in their layer edges, hence, if they are not compatible then some edges connecting two consecutive layers U_i and U_{i+1} cross, where $1 \leq i \leq k-1$. Since these conflicting edges by definition cannot cross any layer one may construct a layer tree which simultaneously is compatible to both original trees as follows: Consider the first layer tree and for an edge from U_i to U_{i+1} which causes a crossing keep the endpoint on U_i but connect it to the endpoint of the crossing edge on U_{i+1}. By construction there is at most one crossing between every two consecutive layers, thus we can simultaneously remove every crossing in a single transformation step, and the resulting tree again is a layer tree.

Theorem 3.7 (Aichholzer et al. [2])**.** *Given a planar point set P with k convex layers. Then the diameter of $\mathcal{T}_{\mathrm{st}}(P)$ is $O(\log k)$.*

Proof. It suffices to show that any crossing-free spanning tree T of P may be transformed into a layer tree by means of at most $O(\log k)$ compatible spanning trees. We will prove this statement constructively.

Consider some triangulation containing T. Then, by Lemma 3.6, for $i \geq 2$ and any point $p \in U_i$ there is an edge pq of the triangulation to a point q on a layer U_j, with $j < i$, such that pq does not cross U_i. Construct a new crossing-free spanning tree T' by choosing these edges for all points in $P \setminus U_1$, and in addition we take all but one edges of U_1. Indeed, T' is connected and does not contain a cycle, and since T' is a subgraph of the triangulation it is also crossing-free. In particular, this also implies that T and T' are compatible.

Note that by construction of T', for any point $p \in P$ there is a path $g_1(p)$ in T' connecting p to U_1 such that only points on layers with strictly decreasing index are visited. Obviously, the union of all these paths is crossing-free. We want to define a similar notion for paths connecting points to the last layer U_k by visiting points only on layers with strictly increasing index such that the union of all paths is crossing-free. Here we have to be a little more careful as we do not have an analog statement of Lemma 3.6 at hand.

For a point $p \in P \setminus U_1$ we define a path $g_k(p)$ from p to the layer U_k which we obtain by the following procedure: If p is not a leaf of the tree T' we follow its edges in an arbitrary manner, however, visiting only points with increasing layer index. Ultimately, we either reach U_k in which case we found the path $g_k(p)$ or we end up in a leaf v of T'. In the latter case consider the shortest segment s_v in $\mathbb{R}^2$ from v to a point on U_k, ties are broken arbitrarily. If the segment does not cross an edge of T' we add this edge to the connection from p to v and define the resulting path to be $g_k(p)$. Otherwise we consider the first edge of T' that is crossed by s_v when walking from v to U_k. This first edge cannot be a layer edge. We slightly rotate s_v around v in direction towards the endpoint of the crossed edge which lies on the layer of larger index until we encounter another point in P. This is similar to the construction shown in Figure 1.28 for the crossing-free partitioning into three classes. Finally, the resulting edge incident to v does not cross T' by construction, and we add it to the previously constructed path and continue recursively.

By the definition of T' any path $g_k(p)$ visits points only on layers of strictly increasing index. Given the construction above it is immediate that $g_k(p)$ neither crosses an edge of T' nor a layer edge of U_k. We claim and require that these paths are also pairwise crossing-free. In order to see this, assume there is p and q such that $g_k(p)$ properly intersects some edge of $g_k(q)$. None of the two edges involved in such a crossing belong to T', hence they both were obtained from rotating the shortest segment from some leaves v_p and v_q of T' to respective points w_p and w_q on U_k. We note that this implies $\|w_p - v_p\| \leq \|w_q - v_p\|$ and $\|w_q - v_q\| \leq \|w_p - v_q\|$. By construction these shortest segments also need to intersect, and we obtain a quadrilateral v_p, v_q, w_p, w_q for which

$$\|w_p - v_p\| + \|w_q - v_q\| \leq \|w_q - v_p\| + \|w_p - v_q\|,$$

i.e., the sum of the diagonals' lengths is at most the sum of lengths of two opposites sides. This, however, contradicts the triangle inequality.

Now, we construct a new graph G on P where, for all points p on layers $U_1, U_2, \ldots, U_{\lceil k/2 \rceil}$ we take the path $g_1(p)$, for the points p on layers $U_{\lceil k/2 \rceil}, \ldots, U_k$ we take the paths $g_k(p)$, and finally, we add all layer edges of U_1 and U_k. By the previous arguments G is a crossing-free graph on P, and it is easily seen to be connected. Furthermore, no edge in G crosses T' nor a layer edge of $U_{\lceil k/2 \rceil}$.

From G we select a spanning tree T'' which contains all but one layer edges of U_1. Then T' and T'' are compatible, and in one more transformation step we may obtain another spanning tree from T'' which contains, in addition, all but one layer edges of both $U_{\lceil k/2 \rceil}$ and U_k. In particular, for small k when $\lceil k/2 \rceil = 2$, or $\lceil k/2 \rceil = k - 1$, and the corresponding layers are consecutive, we indeed find a layer tree with only one edge connecting U_1 and U_2.

Therefore, after three transformation steps we arrive at two independent subproblems each of size at most $\lceil k/2 \rceil + 1$. Therefore, at most $O(\log k)$ steps are necessary to transform T into a layer tree. □

We would like to point out that in the original proof [2] the definition of $g_k(p)$ is slightly different and in fact the corresponding paths are not the same as ours. We favored our description since we used the concepts of rotating segments earlier in the discussion of crossing-free partitions, and also since the condition of a path not crossing T' would have to be made more precise. There are cases where the shortest path from a point p to U_k

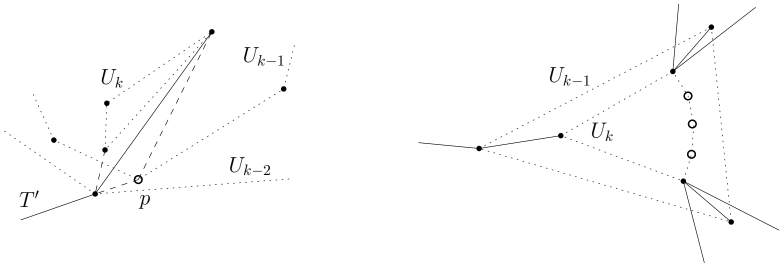

Figure 3.6: Problematic cases for the original proof of Theorem 3.7

which does not cross an edge of T' may actually visit points on layers of smaller index, as shown in the left illustration of Figure 3.6. Here, dotted lines represent layer edges, solid lines indicate edges of T' and dashed edges show paths from p to U_k.

Moreover, it is important to add the layer edges of U_k in the construction above because otherwise G may be disconnected, as can be seen from the right drawing in Figure 3.6. Here, isolated vertices are drawn as circles, layer edges as dotted and paths g_k as solid lines. In order for adding these layer edges to be possible one has to make sure that no edge of T' crosses those layer edges which is achieved by defining T' according to the full power of Lemma 3.6.

Think of me as Yoda.
Only instead of being little and green
I wear suits and I'm awesome.
I'm your bro – I'm Broda!

Barney Stinson,
on "How I met your Mother"

4

Compatible Perfect Matchings

For a planar set of n points in general position we consider the transformation graph whose vertices are the crossing-free perfect matchings of the given point set, and two such perfect matchings are adjacent if their union is also crossing-free.

It was shown by Houle et al. [38] that for any set of points this transformation graph is connected and hence the task of computing its diameter is well-defined. The same authors were able to provide an upper bound of $n-2$ for the diameter, which recently, due to Aichholzer et al. [6], was improved to $O(\log n)$ constituting the current state of the art.

We propose a construction for a set of n points in general position in the plane whose corresponding transformation graph of perfect matchings has diameter $\Omega(\log n / \log \log n)$, nearly providing a tight result on the asymptotic behavior of the upper bound.

Previously to our contribution in [63] only constant lower bounds were known.

4.1 Introduction

Given a set P of n, evenly many, points in the plane let $\mathcal{P}m(P)$ denote the set of all crossing-free straight-line perfect matchings of P. A straight-line embedded graph is called *crossing-free* if every pair of its edges does not share any point other than common endpoints. Two crossing-free perfect matchings M_1 and M_2 of P are *compatible* if their union, i.e., the graph on P with edge set $M_1 \cup M_2$, is crossing-free.

We are interested in the transformation graph $\mathcal{T}_{\mathrm{pm}}(P)$ defined on the vertex set $\mathcal{P}m(P)$ and with edges between compatible perfect matchings. Houle et al. [38] showed that for any set of n points this graph is connected and has diameter at most $n-2$. Recently, Aichholzer et al. [6] improved this upper bound to $O(\log n)$, with at that time no example known for which the diameter is not constant. We give a sub-logarithmic but rather tight lower bound of $\Omega(\log n/ \log\log n)$. We do this constructively by providing point sets of increasing size, and on each point set we specify two perfect matchings achieving the bound.

Although our lower bound construction for the transformation of perfect matchings uses similar ideas as the construction for spanning trees presented in the previous chapter note the problems' immanent difference of connectivity: Spanning trees are connected graphs whereas perfect matchings consist of $n/2$ components. It is this property which makes the construction presented here more challenging.

4.2 The lower bound

In this section we construct planar point sets on which we specify pairs of perfect matchings which need a large number of steps to transform into each other via compatible perfect matchings, i.e., their distance in the transformation graph is large.

We start by introducing the concept of a *prisoner* which is a point serving as a witness that two perfect matchings have at least some fixed distance in $\mathcal{T}_{\mathrm{pm}}(P)$. Based on this we present a recursive construction in order to obtain point sets of increasing size for which the diameter of $\mathcal{T}_{\mathrm{pm}}(P)$ grows as well. This diameter is lower bounded by $\Omega(\log n/ \log\log n)$, where n is the size of the underlying point set. For the sake of a simple description we define point sets with more than two points on a line, i.e., the point sets are not in general position. However, they can easily be changed to do so by applying a small perturbation without losing any of the construction's relevant properties. This is achieved similarly as for the spanning trees.

The key idea is to consider two perfect matchings with a large number of crossings, in particular we will use a first matching with nearly horizontal edges and a second matching with nearly vertical edges. As intuitive as this approach may seem, having many crossings is not enough as can easily be seen by considering at least eight points equidistantly distributed on a circle. Then the following two perfect matchings, one consisting of horizontal edges only and the other one of vertical edges only, have a large number of crossings. However, the diameter of any such transformation graph is 2, since a perfect matching containing only edges on the boundary of the convex hull is adjacent to every other matching.

In order to deal with this issue we will impose dependencies onto the (nearly) horizontal edges such that whatever transformation is made certain (nearly) horizontal edges remain in the matching obtained thereafter.

Before making this statement precise we introduce two further notions. The *granularity* of a point set P is the smallest positive difference of x-coordinates among any two points in P. Recall that a vertical strip R is a point set $R = [\alpha, \beta] \times \mathbb{R}$, for some $\alpha \leq \beta$, its width is $\beta - \alpha$, and an edge blocks R if its endpoints lie on different sides or on the boundary of R.

Given two crossing-free perfect matchings M_1 and M_2 on a point set P, let $p \in P$ and consider the arrangement given by the set of edges in $M_1 \cup M_2$ which are not incident to p. Then p is called *prisoner with respect to M_1 and M_2* if there is a cell C of this arrangement such that $\overline{C} \cap P = \{p\}$, where $\overline{C}$ denotes the closure of the set C.

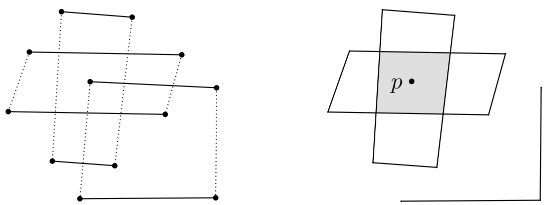

(a) M_1 and M_2 (b) Prisoner p

Figure 4.1: M_1 and M_2 have distance 3 in $\mathcal{T}_{\text{pm}}(P)$

A prisoner guarantees a certain distance of the corresponding perfect matchings M_1 and M_2 in the transformation graph. Furthermore, as we will see later, given a fixed $d \in \mathbb{N}$ we can construct point sets such that the current matching and the target matching define at least one prisoner after any d transformations.

Lemma 4.1. *Let p be a prisoner with respect to M_1 and M_2. Then at least three steps are necessary to transform M_1 into M_2 by compatible matchings.*

Proof. Observe that the existence of a cell C in the arrangement with $\overline{C} \cap P = \{p\}$ implies both $M_1 \neq M_2$, and M_1 and M_2 are not compatible. Hence, their distance in $\mathcal{T}_{\text{pm}}(P)$ is at least 2. Assume it is 2 then there is a perfect matching M compatible to both M_1 and M_2. However, this contradicts $\overline{C} \cap P = \{p\}$ since M matches p to some point outside $\overline{C}$. □

4.2.1 A recursive construction

In the following we describe a way to construct point sets and two perfect matchings whose distance in the transformation graph can be made arbitrarily large. For this purpose consider the point set shown in Figure 4.2(a) given by three copies of a so-called *base gadget* together with a perfect matching consisting of horizontal edges only. The point set has granularity 1/2 assuming a proper coordinate system such that the vertical strip indicated by dashed lines is $[0, 1] \times \mathbb{R}$.

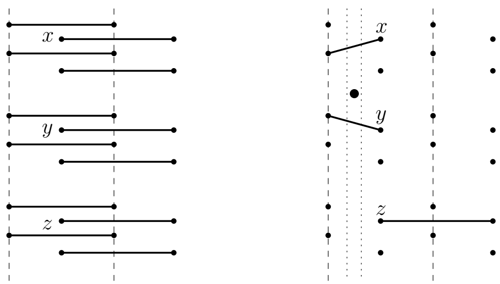

(a) Three copies of the base gadget

(b) After the first transformation

Figure 4.2: Creating a prisoner

After one transformation step to a new compatible perfect matching, the edges incident to the points x, y, and z block vertical strips of width at least 1/2, see Figure 4.2(b). Hence, by the pigeon-hole principle one of the vertical strips $[0, 1/2] \times \mathbb{R}$ or $[1/2, 1] \times \mathbb{R}$ is blocked twice. Now, placing an additional point in-between the two blocking edges creates a candidate for a prisoner. We only need the vertical edges of a second matching indicated by dotted lines for inducing a corresponding cell. We define this second matching at the end of the discussion.

Since there are many ways to transform the initial perfect matching, we have to make sure that for every possible pair of edges that block the same strip there is a candidate prisoner in-between. It is crucial that these new points all have distinct x-coordinates since otherwise we might not obtain prisoners by adding the vertical edges of the second matching. In the example of Figure 4.2(b) this would be the case if the edges incident to x and z after one transformation both block the same strip but the edge incident to y blocks the other. To avoid this we place points in-between the base gadgets so that they equidistantly subdivide each of the vertical strips of width 1/2 into three smaller strips, see Figure 4.3(a).

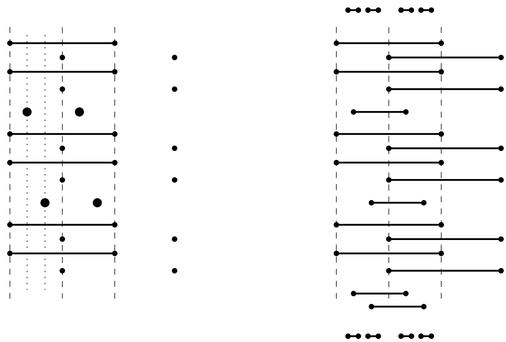

(a) Adding all candidate prisoners

(b) The horizontal matching

Figure 4.3: Point set containing all prisoners

In order to define the second perfect matching of vertical edges we introduce a matching partner for every candidate prisoner with the same respective x-coordinate and place it below the so far constructed point set. We call the hereby obtained set A which will be used in the recursive construction.

Moreover, every candidate prisoner still lacks the two already mentioned vertical edges we need for a prisoner's cell. The points for these edges are placed at the very top and the very bottom of A. Figure 4.3(b) shows the point set A together with the top- and bottom-most points and also the first horizontal matching. Figure 4.4(a) shows the second vertical matching.

Observe that in any matching M which is compatible to the horizontal matching, there is a prisoner with respect to M and the vertical matching, see for instance Figure 4.4(b). Hence, by Lemma 4.1 the diameter of the corresponding transformation graph is at least 4.

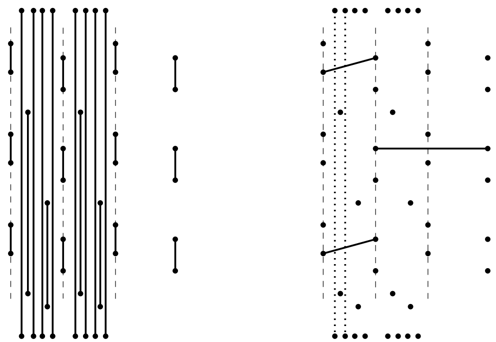

(a) The vertical matching

(b) A prisoner after 1 transformation

Figure 4.4: Point set achieving diameter 4

In order to recursively continue with this idea we need to construct a point set with a prisoner after two transformation steps. For the sake of a comprehensible presentation and since we also perturb the points in the end to guarantee general position we refrain from specifying concrete coordinates for the constructed point sets and rather focus on explaining the construction more carefully.

Recall that we placed the candidate prisoners in such a way that only their matching partners have the same x-coordinate. Hence, a prisoner after one transformation cannot connect to its partner in the next step. By construction the point set A has granularity $1/6$. Therefore, after two transformation steps there is a vertical strip of width at least $1/6$ which is blocked. This is because we placed the candidate prisoners equidistantly in-between the base gadgets.

Now, we vertically stack seven copies of A. By the pigeon-hole principle we know that at least one strip $[(i-1)/6, i/6] \times \mathbb{R}$ is blocked twice after two transformations, for some $1 \leq i \leq 6$. Hence, we are left with defining the new candidate prisoners and their corresponding matching partners. We have to ensure that for every possible pair of blocking edges after two transformation steps there is a single candidate prisoner in-between the blocking edges. This is achieved in the following way. For $1 \leq i \leq 6$, we separately consider the strip $[(i-1)/6, i/6] \times \mathbb{R}$ in which we place six points, one in-between consecutive copies of A, and equidistantly distributed inside $[(i-1)/6, i/6]$.

Hence, in total we add 36 candidate prisoners and equally many matching partners. Note that this new point set, call it B, has granularity $1/(6 \cdot 7) = 1/42$. Similarly to Figure 4.4(a), we add top- and bottom-most points for each candidate prisoner which we need for defining the prisoner's cell. Recall once again that these points, however, are not part of the recursive construction as they do not belong to B.

Starting with the horizontal matching, after any two compatible transformation steps we obtain a new perfect matching which together with the vertical perfect matching induces a prisoner. Hence, the corresponding transformation graph has diameter at least 5.

By the same argument, stacking 43 copies of B and spreading in all candidate prisoners we obtain a new point set C which in turn, by adding top- and bottom-most points as before, yields a transformation graph with diameter at least 6. We observe that the granularity of C has already decreased to $1/(42 \cdot 43) = 1/1806$.

We omit the exact calculation of the diameter's asymptotic behavior in terms of the number of points used in this construction because we will drastically improve on it in the following section. However, note the doubly exponential decrease of the granularity and accordingly the doubly exponential growth of the number of previously constructed point sets used in the recursion. Without proof we mention that the construction leads to point sets of size n with perfect matchings M_1 and M_2 such that if d steps are needed to transform M_1 into M_2 then $n = O\big(2^{2^d}\big)$, that is $d = \Omega(\log\log n)$.

4.2.2 Many prisoners help

In the following we will further develop the concept for constructing point sets whose transformation graphs have large diameter. Recall that the sufficient condition for applying the pigeon-hole principle in the recursion is that the number of copies of previously constructed point sets is strictly larger than the inverse of the granularity.

We will now subdivide the vertical strips by even more candidate prisoners, still equidistantly, in order to reduce the number of copies we need such that the diameter of the transformation graph increases by 1. This is motivated by the following observation when stacking many copies. We did not yet take into account that (a lot) more than just one strip may be blocked (a lot) more often than just twice. Still every candidate prisoner in-between two blocking edges cannot connect to its partner after the next transformation step and hence will be incident to an edge blocking a strip whose width is at least the granularity of the point set.

In order to make this statement precise we consider a single strip that is blocked at least twice and analyze what happens after the next transformation. For every pair of consecutive blocking edges we appoint the candidate prisoner in-between with the largest y-coordinate to be *responsible* for the small vertical strip to its left.

We show in the following that there is an injective map from the set of responsible candidate prisoners to the set of edges blocking the smaller strips after the next transformation step.

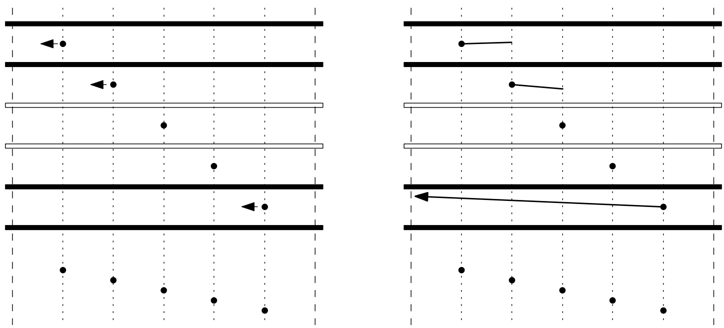

(a) Responsible points (b) Many blocked strips

Figure 4.5: What happens in one large strip

In the example depicted in Figure 4.5(a) we see the magnification of a single strip indicated by dashed vertical lines to the left and right end. The six copies of the previous recursive construction step are shown as rectangles. They are drawn empty if there is no edge blocking the strip, and solid if the strip is blocked. The five prisoners between the copies subdivide the strip into six smaller strips indicated by vertical dotted lines. The arrows indicate the prisoners responsible for their corresponding strip to the left.

After the next transformation step none of the responsible prisoners connects to its matching partner. Thus, considering the granularity of the point set, their incident edges either block the smaller strip to their left or to their right, see Figure 4.5(b). Note that the candidates are responsible for their *left* smaller strip. If they connect to the right then their left strip remains available unless a candidate with larger y-coordinate already claimed it. By construction, once a responsible prisoner connects to the left its incident edge has to block *all* smaller strips until the very left end, and at least one of these small strips was not yet accounted for.

In particular this implies that every edge which blocks the initial strip, except for the top-most, guarantees a blocked strip of width equal to the granularity after the next transformation.

We are now ready to explain the more advanced recursive construction using this observation. Denote the base gadget from the previous subsection by A_0 and, for $k \geq 1$, let $a_k \in \mathbb{N}$ be the number of copies of A_{k-1} we use to construct A_k. Denote $\pi_k = \prod_{i=1}^{k} a_i$, for $k \geq 1$, and notice that the number of strips in A_k is $2\pi_k$ because each strip of A_{k-1} is subdivided into a_k smaller strips, and A_0 has two strips. As before, from A_k one can easily construct a point set and define the corresponding horizontal and vertical matchings such that there exists a prisoner after k transformations.

Let $b_k \in \mathbb{N}$ denote the minimum number of strips of width $(2\pi_k)^{-1}$ that are blocked in A_k after any k transformation steps. For the construction to work we need to be able to apply the pigeon-hole principle in each recursion step. A necessary and sufficient condition for a point set to yield a blocker containing A_d after d transformations is $b_k > 0$, for all $1 \leq k \leq d$.

Lemma 4.2. *For $k \geq 1$, the integer sequence b_k satisfies*

$$b_k = \pi_k \left(1 - 2\sum_{i=1}^{k} a_i^{-1}\right).$$

Proof. We shall give a proof by induction on k. To this end recall that in A_0 there are two strips, one of which is blocked after the first step. Using the observation about responsible prisoners, in A_1 there are $2a_1 = 2\pi_1$ strips, of which at least $a_1 - 2 = \pi_1(1 - 2/a_1)$ are blocked after two steps.

For the inductive step assume that at least b_k strips of width $(2\pi_k)^{-1}$ are blocked in A_k after k steps, and $b_k > 0$. Then stacking a_{k+1} copies of A_k yields after the next transformation at least

$$\begin{aligned} b_{k+1} &= a_{k+1} b_k - 2\pi_k \qquad (4.1) \\ &= \pi_{k+1}\left(1 - 2\sum_{i=1}^{k} a_i^{-1}\right) - 2\pi_k = \pi_{k+1}\left(1 - 2\sum_{i=1}^{k+1} a_i^{-1}\right) \end{aligned}$$

blocked strips of width $(2\pi_{k+1})^{-1}$. □

Since $a_k > 0$, for all k, the condition $b_d > 0$ for our target blocker is equivalent to $b_{d-1} a_d > 2\pi_{d-1}$ by identity (4.1) which immediately implies $b_{d-1} > 0$. Recursively, $b_k > 0$ follows for all remaining k. Thus, $b_d > 0$

and, therefore, by Lemma 4.2

$$2\sum_{i=1}^{d} a_i^{-1} < 1, \tag{4.2}$$

is a necessary and sufficient condition for the blocking property of A_d.

Now, we turn to the number of points used in the construction. In addition to the copies of previously constructed point sets we also spread candidate prisoners in-between. Let $n_k := |A_k|$, for $k \geq 0$, where $n_0 = 8$.

Lemma 4.3. *For $k \geq 1$, the number of points in A_k is*

$$n_k = \pi_k \left(n_0 + 4k - 4\sum_{i=1}^{k} a_i^{-1} \right).$$

Proof. By construction we have $n_1 = a_1 \cdot n_0 + 2(a_1 - 1) = \pi_1(n_0 + 2 - 2/a_1)$ as induction base. Moreover, for constructing A_{k+1} we stack a_{k+1} copies of A_k and add the candidate prisoners and their matching partners, thus

$$n_{k+1} = a_{k+1} n_k + 2 \cdot 2\pi_k(a_{k+1} - 1),$$

which implies the statement of the lemma after substituting the inductive hypothesis for n_k. □

With $a_k = 2d + 1$, for all $1 \leq k \leq d$, condition (4.2) is fulfilled since

$$2\sum_{i=1}^{d} a_i^{-1} = \frac{2d}{2d+1} < 1.$$

By Lemma 4.3 this choice yields a set A_d of size

$$n_d = (2d+1)^d \left(n_0 + 4d - \frac{4d}{2d+1} \right) = O\big(c^d d^d\big),$$

for some constant $c > 0$. Adding top- and bottom-most points needed for the vertical edges of the prisoners' cells we obtain a point set with at most $2n_d$ points for which the diameter of the transformation graph is $d + 3$.

Theorem 4.4. *For arbitrarily large n there is a set P of n points in the plane for which the diameter of $\mathcal{T}_{\mathrm{pm}}(P)$ is $\Omega(\log n / \log\log n)$.*

The choice of the a_k is asymptotically optimal in our construction here, as (4.2) and Lemma 4.3 imply $n_d > \pi_d$, and the geometric-harmonic means inequality yields

$$\sqrt[d]{\pi_d} \geq \frac{d}{\sum_{i=1}^{d} a_i^{-1}} > 2d.$$

Part III

Counting Crossing-Free Geometric Graphs

Chuck Norris counted to infinity.
Twice!

Chuck Norris Facts

5

Counting Crossing-Free Graphs with Exponential Speed-Up

We show that one may count the number of crossing-free geometric graphs on a given planar point set exponentially faster than enumerating them. More precisely, given a set P of n points in general position in the plane we can compute $\mathsf{pg}(P)$, the number of plane graphs on P, in time at most $\frac{\mathrm{poly}(n)}{\sqrt{8}^n} \cdot \mathsf{pg}(P)$. There are no similar statements known for other graph classes like triangulations, spanning trees or perfect matchings.

The exponential speed-up is obtained by enumerating the set of triangulations on P and then, without repetition, for each triangulation count all its subgraphs. For a set P of n points with triangular convex hull we further improve the base of the exponential from $\sqrt{8} \approx 2.828$ to 3.347. As a main ingredient for this refinement we show that there is a constant $\alpha > 0$ such that a triangulation on P chosen uniformly at random contains, in expectation, at least n/α non-flippable edges. The best value for α we obtain is 37/18.

This is joint work with Emo Welzl [66].

5.1 Preliminaries

Let P be a finite set of points in the plane. We assume that P is in general position, i.e., no three points are collinear and no four points cocircular. A *geometric graph on P* is a simple graph defined on the vertex set P whose edges are straight segments connecting the corresponding endpoints. Such a straight-line embedded graph is *crossing-free* if no pair of its edges shares any point except for, possibly, a common endpoint.

A crossing-free graph which is maximal with respect to the number of edges is called *triangulation*. By definition a graph consisting of only one vertex without any edge, or two vertices joined by a single edge also constitute triangulations. However, for the sake of keeping the intuition about triangular faces and a closed presentation of the results we only consider sets consisting of at least three points.

In the following we assume that the underlying point set P is fixed and write $n := n(P) = |P|$ for its cardinality. Furthermore, let $k := k(P)$ denote the number of points on the boundary of the convex hull of P, thus by assumption $n \geq k \geq 3$. Recall Euler's polyhedral formula for any plane graph G on P which states that

$$n + f - e = 2,$$

where f denotes the number of faces and e the number of edges in G. In a triangulation every bounded face is a triangle and the unbounded face is of length k, hence we find that $3(f-1) + k = 2e$ because every edge belongs to exactly two faces. Substituting $f - 1 = e - n + 1$ from Euler's formula yields $e = 3n - k - 3$, and in turn for the number of bounded faces we obtain $f - 1 = 2n - k - 2$. Hence, any triangulation consists of exactly

$$M := M(P) = 3n - k - 3 \tag{5.1}$$

edges, which we will further discriminate as follows. An edge in a triangulation T is called *flippable* if it is contained in the boundary of two triangles of T whose union is a convex quadrilateral, see Figures 5.1(a) and (b); Otherwise the edge is called *non-flippable*, see Figure 5.1(c). We write $\mathit{fl}(T)$ for the number of flippable edges in T, and similarly $\mathit{nfl}(T)$ for the number of non-flippable edges. Note that for any triangulation T, clearly, $\mathit{fl}(T) + \mathit{nfl}(T) = M$, and $\mathit{nfl}(T) \geq k$ holds since $k \geq 3$ and the edges on the boundary of the convex hull of P are always non-flippable by definition.

We are interested in the number of plane graphs that can be defined on P which we denote by $\mathsf{pg}(P)$. The currently best upper bound [64] for this

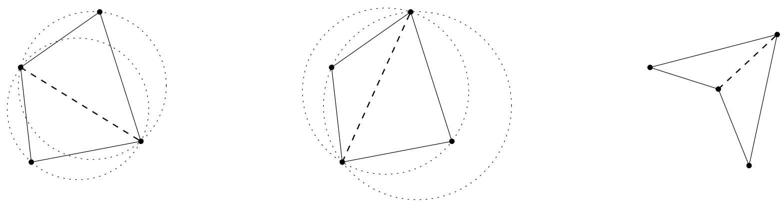

(a) (b) (c)

Figure 5.1: Flippable edges in (a) and (b); Non-flippable edge in (c)

quantity stands at 343.106^n, where n is the number of points in P. This bound is a result jointly with Jack Snoeyink and Emo Welzl which we shall prove in Chapter 6.

The set of all triangulations on P is denoted by $\mathcal{T}r(P)$, and we will write $\mathsf{tr}(P) := |\mathcal{T}r(P)|$ for its cardinality. The best known upper bound for $\mathsf{tr}(P)$ is 43^n due to Sharir and Welzl [72].

Considering lower bounds it was common belief that a point set in general position always allows for exponentially many triangulations, however, surprisingly a rigorous treatment of this question just received attention quite recently. Note that general position is crucial here as we will also see in the next chapter. A first proof may be attributed to Galtier et al. [32] who showed that any triangulation on n points contains at least $\frac{n-4}{6}$ edges that can be flipped simultaneously, although it was Aichholzer et al. [9] who first mentioned its implication towards providing a lower bound of $2^{(n-4)/6} \approx 1.124^n$ triangulations on n points. In the same paper [9] this result was further improved to 2.012^n, the first bound of the form $\Omega\big((2+\beta)^n\big)$ with $\beta > 0$ constant. The current record of $\Omega(2.338^n)$ triangulations any set of n points in general position has is due to Aichholzer et al. [10]. For point sets with k vertices on the boundary of the convex hull McCabe [57] and McCabe and Seidel [58] gave even better bounds, the currently best are $\Omega\left((\frac{30}{11})^k(\frac{11}{5})^{(n-k)}\right)$ and, for any k fixed, $\Omega(2.63^n)$.

In this chapter we will show that there is an absolute constant $c > 1$ such that $\mathsf{pg}(P) \geq c^n \cdot \mathsf{tr}(P)$ for any planar set P of $n \geq 3$ points in general position, while we are still able to compute $\mathsf{pg}(P)$ in time necessary to enumerate $\mathcal{T}r(P)$ times a small polynomial factor in n. The best value for the constant c we obtain is $\sqrt{8}$. Such an enumeration for the set of triangulations is achieved by applying the reverse search technique devised by Avis and Fukuda [15]. The fastest algorithm for this enumeration is due to Bespamyatnikh [20] and needs time $O(\log\log n)$ per output triangulation.

Recently, and independently of our work, Katoh and Tanigawa [45] suggested an idea for enumerating crossing-free geometric graph classes relatively similar to our approach by introducing a lexicographic order on the set of triangulations.

To proceed with the presentation, a flippable edge in T is called *Lawson edge* if the circumcircle of each boundary triangle also contains the respective other boundary triangle in its interior. See for instance the dashed line segment from Figure 5.1(b) in contrast to Figure 5.1(a). Observe that this notion is well-defined since we assume general position. We denote by $L(T)$ the set of Lawson edges in T, and by $\ell(T)$ its cardinality. A triangulation T for which $L(T) = \emptyset$ is called *Delaunay triangulation of* P and it is well-known that such a triangulation exists on any point set P and it is unique if P is in general position. Lawson [51] showed that, when starting with any triangulation on n points, the process of repeatedly flipping a Lawson edge in the current triangulation terminates with the Delaunay triangulation after $O(n^2)$ flips.

We require a more general setting, for this purpose we recall a few definitions and facts about the constrained Delaunay triangulation which was introduced by Lee and Lin [52]. Given a crossing-free geometric graph G on P, then two points $p, q \in P$ are *visible* from each other if the line segment pq does not intersect the interior of any edge in G. The *constrained Delaunay triangulation* $T^*(G)$ of G is a triangulation containing the edges of G which has the additional property that the circumcircle of each triangle in $T^*(G)$ does not contain any other point which is visible from all three vertices of the triangle. Lee and Lin showed that for any graph G the constrained Delaunay triangulation $T^*(G)$ exists and is unique if P is in general position. Moreover, they proved that $T^*(G)$ is obtained from any triangulation containing $E(G)$ by repeatedly flipping a Lawson edge distinct from $E(G)$ as long as applicable. We shall give a proof of these facts for the sake of being self-contained.

Theorem 5.1. *For any crossing-free graph G on a point set P in general position the constrained Delaunay triangulation $T^* = T^*(G)$ of G exists and is unique. Moreover, it is characterized by $L(T^*) \subseteq E(G) \subseteq E(T^*)$.*

Proof. For a triangulation T containing the edges of G we consider the vector $\chi(T) \in \mathbb{R}^{2n-k-2}$ that stores the smallest angle of each bounded triangle of T in non-decreasing order. We note that in this lexicographic order flipping a Lawson edge strictly increases the vector χ.

Indeed, let cd in T be a Lawson edge and ab the flipped edge in T', see Figure 5.2. Consider the entries of the vectors $\chi(T)$ and $\chi(T')$ which

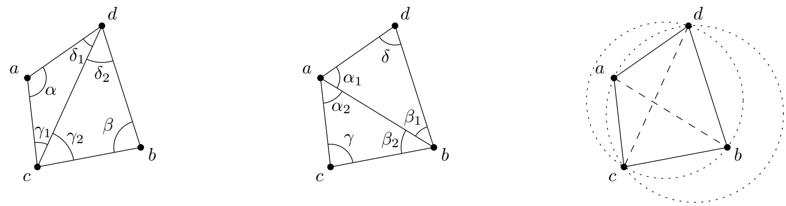

Figure 5.2: The angles before and after flipping a Lawson edge

change during the flip, that is where the entries

$$\min\{\alpha, \gamma_1, \delta_1\}, \min\{\beta, \delta_2, \gamma_2\} \quad \text{become} \quad \min\{\delta, \alpha_1, \beta_1\}, \min\{\gamma, \beta_2, \alpha_2\}.$$

All angles are strictly positive because of general position, hence with cd being a Lawson edge it follows from the Inscribed Angle Theorem that

$$\begin{array}{lll} \delta > \delta_1, \delta_2 & \alpha_1 > \gamma_2 & \alpha_2 > \delta_2 \\ \gamma > \gamma_1, \gamma_2 & \beta_1 > \gamma_1 & \beta_2 > \delta_1, \end{array}$$

and in particular

$$\min\{\delta, \alpha_1, \beta_1, \gamma, \beta_2, \alpha_2\} > \min\{\alpha, \gamma_1, \delta_1, \beta, \delta_2, \gamma_2\}.$$

Therefore, $\chi(T') > \chi(T)$ and repeatedly flipping Lawson edges not belonging to $E(G)$ must terminate eventually since there is only a finite number of triangulations on P. Let T^* be the resulting triangulation of this process. Clearly, $L(T^*) \subseteq E(G) \subseteq E(T^*)$ by construction. It remains to show that T^* is a constrained Delaunay triangulation, i.e., satisfies the visibility criterion, and it is the unique triangulation with that property.

Suppose for contradiction that T^* is not a constrained Delaunay triangulation, then there is a witness triangle in T^* whose circumcircle contains a point which is visible from all three vertices of the triangle. Let p be this point and abc the witness triangle such that the line through bc separates a and p. Note that because of the visibility constraint no edge crossing the segment ap can be an edge of G.

Now, consider the set Δ of all triangles in T^* intersecting the segment ap, see the left image of Figure 5.3. Assume this finite set to be ordered according to the intersection with ap, the first triangle being $\Delta_1 = abc$ and the last one Δ_j containing p, i.e., $\Delta = (\Delta_1, \Delta_2, \ldots, \Delta_j)$. We take a closer look at the second triangle $\Delta_2 = cbd$ of this ordered set. If the point d happens to lie inside the circumcircle of abc then by definition the edge bc

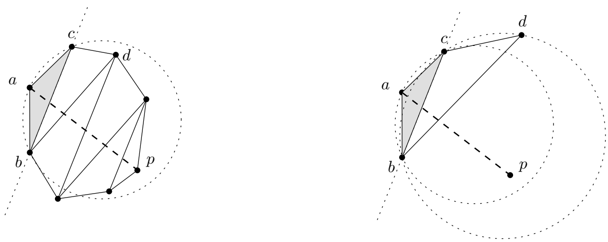

Figure 5.3: Witness triangle for the constrained Delaunay triangulation

is a Lawson edge of T^*. However, since bc crosses ap it cannot be an edge of G by the visibility constraint, which contradicts the fact that no edge T^* can be flipped. So we may assume that d lies outside the circumcircle of abc. In this case, however, the circumcircle of the second triangle cbd still contains p. In order to see this recall that the line through bc divides the circumcircle of abc into two parts one of which contains p. This part clearly is contained inside a circle that goes through the points c, b and another point outside, see the right illustration of Figure 5.3.

Hence, we may now continue with the argument along the triangles $\Delta_3, \ldots, \Delta_j$. For $3 \leq i \leq j$, either Δ_i is contained in the circumcircle of Δ_{i-1}, in which case we found a Lawson edge, or the circumcircle of Δ_i still contains p and we continue with Δ_{i+1}. At some point, which happens the latest when considering the last triangle Δ_j containing p, we find an edge which does not belong to $E(G)$ but is a Lawson edge for T^*. This gives the desired contradiction.

In order to show that T^* is unique, and in particular does not depend on the triangulation with which we start the flipping process, we show that T^* is the largest triangulation with respect to the lexicographic order of χ. For this purpose assume for contradiction that $T' \neq T^*$ is the largest triangulation containing $E(G)$ and $\chi(T') \geq \chi(T^*)$. By assumption no edge may be flipped in T' which yields $L(T') \subseteq E(G) \subseteq E(T')$, and by the previous argument T' has to satisfy the visibility condition and hence is an other constrained Delaunay triangulation.

Since $T' \neq T^*$ there are edges $e \in E(T^*)$ and $f \in E(T')$ that cross. In particular we may choose $e = bc$ and $f = ad$ such that abc forms a triangle in T^*, and there is no further intersection of an edge in T' with e which is closer to b, see Figure 5.4. We observe that the last property guarantees that abd is a triangle in T'. Neither e nor f belongs to $E(G)$, thus d is visible from a and c is visible from b. Due to general position one

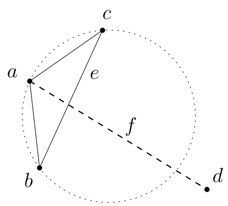

Figure 5.4: Choosing crossing edges e and f in T^* and T', respectively

of the circumcircles of abc or abd contains the respective fourth point, and since T' and T^* are both constrained Delaunay triangulations the points c and d cannot be visible from each other. Hence, there exists another point of P inside the triangle given by the intersection of e with f and the vertices d and c. The closest such point to e is visible from a, b, c, and correspondingly the closest point to f is visible from a, b, d. At least one of them is contained inside the respective circumcircle of abc or abd, implying that T' or T^* does not satisfy the visibility constraints. This is a contradiction. □

5.2 Extraction from triangulations

The following theorem is the key to counting and estimating the number of crossing-free geometric graphs in terms of the number of triangulations on a point set P. The basic ingredient is to partition the set of all crossing-free graphs by associating each graph with its constrained Delaunay triangulation. Then the theorem suggests an algorithm for computing $\mathsf{pg}(P)$ in time $O\big(\text{poly}(n) \cdot \mathsf{tr}(P)\big)$ by enumerating $\mathcal{T}r(P)$ according to the reverse search method due to Avis and Fukuda [15] and counting all subgraphs associated with the same triangulation. We will also show that there is a constant $c > 1$ such that $\mathsf{pg}(P) \geq c^n \cdot \mathsf{tr}(P)$, implying that one may count $\mathsf{pg}(P)$ exponentially faster than enumerating all graphs.

Theorem 5.2. *Given a set P of $n \geq 3$ points in general position in the plane let M denote the number of edges in any triangulation on P, and for a fixed triangulation T we write $\ell(T)$ for the number of its Lawson edges. Then*

$$\mathsf{pg}(P) = \sum_{T \in \mathcal{T}r(P)} 2^{M-\ell(T)}. \tag{5.2}$$

Proof. Consider the following partition of the set of crossing-free geometric graphs on P. For every triangulation T on P there is a partition class consisting of all crossing-free subgraphs G of T that contain the set of Lawson edges of T, i.e., for which

$$L(T) \subseteq E(G) \subseteq E(T). \tag{5.3}$$

Indeed, this defines a partition due to the existence and uniqueness of the constrained Delaunay triangulation from Theorem 5.1. The partition class associated with a triangulation T contains exactly $2^{M-\ell(T)}$ crossing-free graphs. Summing over all triangulations yields the statement. □

We consider the set P_6, depicted in Figure 5.5, consisting of six points in general position. It can easily be checked that there are exactly six triangulations on P_6, the corresponding Lawson edges are drawn as dashed line segments. Indeed, notice that any triangulation of P_6 has to contain the four boundary edges as well as two edges connecting each of the interior points with two vertices on the boundary. No edge may cross any of these eight edges. Now, either the interior vertices are adjacent in which case there are four triangulations since there are two possibilities to triangulate each of the convex quadrangles, or a pair of opposite boundary vertices is adjacent in which case the triangulations are unique. With four points on the convex hull every triangulation contains $M = 11$ edges, hence by Theorem 5.2 we immediately obtain

$$\mathsf{pg}(P_6) = 2^{11} + 2^{10} + 2^{10} + 2^{10} + 2^{9} + 2^{10} = 6656.$$

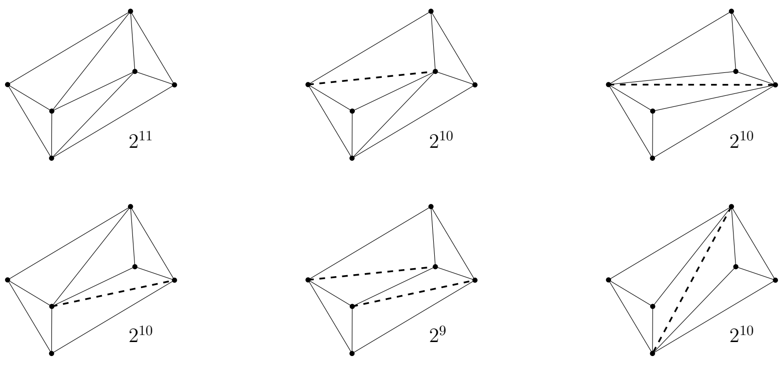

Figure 5.5: Counting crossing-free graphs on a set of six points

Actually, from the proof of Theorem 5.2 one may obtain counting algorithms for any class of crossing-free graphs. We simply iterate over the set of all triangulations on P, and for a triangulation T we count the members G of the desired graph class that fulfill the edge containment property (5.3). Observe that computing $L(T)$, and hence also $\ell(T)$, can be done in linear time. In addition to counting all crossing-free geometric graphs on P we shortly discuss two further prominent examples, counting perfect matchings and counting spanning trees. For certain graph classes there are efficient algorithms to count the number of perfect matchings and spanning trees that are contained as subgraphs. We refer to [41, Chapter 1] for a more detailed description and proofs of the following results.

Perfect matchings. Kasteleyn [43] described the class of *Pfaffian orientable graphs*, for which he found an elegant and very efficient way to count the number of perfect matchings. An orientation of an undirected graph G is called *Pfaffian* if for any two perfect matchings M_1 and M_2 in G the following holds: When walking along the edges of any cycle in $M_1 \cup M_2$ we encounter an odd number of edges with the same orientation as is imposed on G. Note that all cycles in $M_1 \cup M_2$ are even, hence the direction for the traversal of the cycles does not matter. Given a directed graph $\vec{G}$ its adjacency matrix $A(\vec{G})$ is defined by

$$a_{ij} = \begin{cases} +1 & (i,j) \in E(\vec{G}), \\ -1 & (j,i) \in E(\vec{G}), \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 5.3 (Kasteleyn)**.** *For any Pfaffian orientation $\vec{G}$ of G the number of perfect matchings in G is $\sqrt{\det A(\vec{G})}$.*

This notion turns out useful as all planar graphs allow for Pfaffian orientations. Coming back to our counting algorithm, given a triangulation T we want to compute the number of perfect matchings in T containing $L(T)$. Obviously, $L(T)$ has to be a matching otherwise we may safely continue with the next triangulation. If $L(T)$ is a matching we remove all vertices which are incident to an edge of $L(T)$ and consider the subgraph of T spanned by the remaining vertices. Since this graph is planar as well we can compute its number of perfect matchings by Theorem 5.3.

One may observe that for a fixed T the calculation above can be done in time polynomial in the number of vertices. Therefore, we obtain an algorithm for counting the total number of perfect matchings on a given set

P of n points in time poly$(n)\cdot\mathsf{tr}(P)$. However, the number of triangulations can be exponentially larger than the number of perfect matchings. For instance we already noted in Chapter 1 that a set Γ_n in convex position, with n even, has $C_{n-2} = \Theta^*(4^n)$ triangulations but only $C_{n/2} = \Theta^*(2^n)$ perfect matchings, where $C_n = \frac{1}{n+1}\binom{2n}{n}$ denotes the n-th Catalan number.

Spanning trees. Regarding the computation of the number of spanning trees in a given planar graph, i.e., a graph which allows for a plane embedding, already Kirchhoff [46] knew about a beautiful identity, often referred to as "Matrix Tree Theorem". Given a graph G let $A = A(G)$ denote its adjacency matrix and $D = D(G)$ the diagonal matrix with the vertices' degrees on its main diagonal. Moreover, we write $(D-A)_{ii}$ for the submatrix of $D - A$ obtained by deleting its i-th row and i-th column.

Theorem 5.4 (Kirchhoff). *Let G be a loop-free graph on n vertices. Then the number of spanning trees in G is* $\det(D-A)_{ii}$, *for any* $1 \leq i \leq n$.

A slightly more general extension which we need is due to Tutte [79] who proved that the statement of Theorem 5.4 also holds if D and A are defined for multigraphs, i.e., the ij-th entry of the adjacency matrix A is the number of edges between vertices i and j, and the degrees are the entries on the main diagonal of D. We also mention a work by Moon [59] for a description of several related results and identities for counting trees.

Given a triangulation T our algorithm should compute the number of spanning trees in T containing $L(T)$. Clearly, necessary for their existence is that $L(T)$ induces a forest on P. In order to guarantee that every spanning tree contains $L(T)$ we contract the corresponding edges in the triangulation and thus obtain a multigraph. Any loops that occur during this procedure may be deleted. Indeed, if we would consider such a loop to be part of a spanning tree it would complete a cycle after undoing the contraction of $L(T)$. Moreover, observe that at any time during the contraction if there are multiple edges between two vertices at most one of them corresponds to an edge in $L(T)$, as otherwise initially $L(T)$ was not cycle-free.

After contracting all edges, and removing possible loops, we obtain a multigraph for which we may compute the number of its spanning trees in polynomial time. This value is exactly the number of spanning trees in T containing all its Lawson edges. Therefore, we have a counting procedure that on input P, a set of n points, outputs the number of crossing-free spanning trees on P in time poly$(n)\cdot\mathsf{tr}(P)$.

We conjecture that every point set allows for more spanning trees than triangulations, and perhaps even exponentially more in which case the algorithm above yields an exponential speed-up. However, this question is open to the best of our knowledge. We mention some related results showing that a straight-forward proof idea, of estimating the number of spanning trees in a triangulation and comparing it to the number of triangulations a fixed spanning tree can be extended to, fails. The first result by Rote et al. [68] states that any plane graph on n vertices has at most $O\big((16/3)^n\big) = O(5.334^n)$ spanning trees, which was recently improved to $O(5.286^n)$ by Buchin and Schulz [23]. On the other hand, consider the double zig-zag chain introduced by Aichholzer et al. [8] which allows for spanning trees that are contained in every of its $\Theta^*(\sqrt{72}^n)$ triangulations, where $\sqrt{72} \approx 8.485$. Indeed, the set of edges appearing in every triangulation forms a connected graph.

All plane graphs. In the remainder of this section we will show that counting all crossing-free geometric graphs by enumerating all triangulations $T \in \mathcal{T}r(P)$ and summing up $2^{M-\ell(T)}$ yields an exponential speed-up compared to the enumeration of all graphs. Dividing both sides of identity (5.2) by the total number of triangulations on P we obtain

$$\frac{\mathsf{pg}(P)}{\mathsf{tr}(P)} = \sum_{T \in \mathcal{T}r(P)} 2^{M-\ell(T)} \cdot \tfrac{1}{\mathsf{tr}(P)} = \mathbb{E}\left[2^{M-\ell(T)}\right], \tag{5.4}$$

where the expectation of the random variable $2^{M-\ell(T)}$ is understood with respect to the uniform distribution over all triangulations on P. In order to estimate the expected value we employ a standard tool from probability theory. Although its statement is true in a much more general setting we restrict ourselves to presenting a version sufficiently strong for our purpose.

Lemma 5.5 (Jensen's Inequality)**.** *Let $\varphi : \mathbb{R} \to \mathbb{R}$ be a convex function and $t_i \in [0,1]$, for $1 \le i \le N$, such that $\sum_{i=1}^N t_i = 1$. Then, for any $x_i \in \mathbb{R}$,*

$$\sum_{i=1}^N \varphi(x_i)t_i \ge \varphi\left(\sum_{i=1}^N x_i t_i\right).$$

Proof. We give a proof by induction on N. If $N = 1$ the statement is trivial since $t_1 = 1$, and for $N = 2$ where $t_1 + t_2 = 1$, the inequality $\varphi(x_1)t_1 + \varphi(x_2)(1-t_1) \ge \varphi(x_1t_1 + x_2(1-t_1))$ is exactly the definition of φ being convex, thus correct.

For the induction step from $N-1$ to N, assume that $t_N < 1$ otherwise we are done immediately. Therefore,

$$\sum_{i=1}^{N} \varphi(x_i)t_i = \left(\sum_{i=1}^{N-1} \varphi(x_i)\frac{t_i}{1-t_N} \right) \cdot (1-t_N) + \varphi(x_N)t_N.$$

Since $\sum_{i=1}^{N-1} \frac{t_i}{1-t_N} = 1$ we may apply the induction hypothesis to find that the right-hand side in the equation above is at least

$$\varphi\left(\sum_{i=1}^{N-1} x_i \frac{t_i}{1-t_N} \right) \cdot (1-t_N) + \varphi(x_N)t_N \geq \varphi\left(\sum_{i=1}^{N} x_i t_i \right),$$

where the last inequality holds due to the convexity of φ. □

Now, we translate Jensen's Inequality to our setting of a random variable X taking values in some finite set $S \subseteq \mathbb{R}$ and $\varphi : \mathbb{R} \to \mathbb{R}$ being a convex function. Then, writing $p_x := \mathbb{P}[X = x]$, for $x \in S$, and noting that $\sum_{x\in S} p_x = 1$, we conclude

$$\mathbb{E}[\varphi(X)] = \sum_{x\in S} \varphi(x)p_x \geq \varphi\left(\sum_{x\in S} x p_x \right) = \varphi(\mathbb{E}[X]).$$

Since the exponential function $x \mapsto 2^x$ is convex and $2^{M-\ell(T)}$ only takes values in the finite set $\{1, 2, 4, \ldots, 2^M\}$ we may apply Jensen's inequality to derive the following lower bound for the expectation in (5.4)

$$\mathbb{E}\left[2^{M-\ell(T)}\right] \geq 2^{\mathbb{E}[M-\ell(T)]} = 2^{M-\mathbb{E}[\ell(T)]},$$

where we also use linearity of expectation. Hence, by providing an upper bound on $\mathbb{E}[\ell(T)]$, i.e., the expected number of Lawson edges in a uniformly at random chosen triangulation on P, we obtain a lower bound for the number of plane graphs on P in terms of the number of triangulations

$$\mathsf{pg}(P) \geq 2^{M-\mathbb{E}[\ell(T)]} \cdot \mathsf{tr}(P). \tag{5.5}$$

We recall that *fl*(T) denotes the number of flippable edges in a triangulation T, and *nfl*(T) stands for the number of non-flippable edges in T.

Lemma 5.6. *For any set P of points in general position it holds that*

$$2 \cdot \mathbb{E}[\ell(T)] = \mathbb{E}[\mathrm{fl}(T)].$$

Proof. Let S be the set of pairs (e, T) with T a triangulation on P and e a flippable edge in T. Consider an element (e, T) of this set and let e' be the other diagonal of the convex quadrilateral consisting of the boundary triangles of e in T. When flipping e in T, i.e., replacing e by e', we obtain a new triangulation T'. Clearly, $(e', T') \in S$ and flipping e' in T' yields T again. Hence, there is a (canonical) perfect matching between the elements of S. Note that by definition either e or e' is a Lawson edge of its respective triangulation. Therefore,

$$\begin{aligned} |S| &= \sum_{T \in \mathcal{T}r(P)} \mathit{fl}(T) = \mathbb{E}[\mathit{fl}(T)] \cdot \mathsf{tr}(P) \\ \frac{|S|}{2} &= \sum_{T \in \mathcal{T}r(P)} \ell(T) = \mathbb{E}[\ell(T)] \cdot \mathsf{tr}(P), \end{aligned}$$

which proves the statement since $\mathsf{tr}(P) \neq 0$. □

Recall that $\mathit{fl}(T) + \mathit{nfl}(T) = M$, which together with Lemma 5.6 gives

$$\mathbb{E}[\ell(T)] = \frac{1}{2} \cdot (M - \mathbb{E}[\mathit{nfl}(T)]),$$

using linearity of expectation. Plugging this into inequality (5.5) we obtain the following estimate.

Theorem 5.7. *For any set P of $n \geq 3$ points in general position in the plane the following holds*

$$\mathsf{pg}(P) \geq 2^{(M + \mathbb{E}[\mathit{nfl}(T)])/2} \cdot \mathsf{tr}(P). \tag{5.6}$$

Recall that for any triangulation T on P we have $\mathit{nfl}(T) \geq k$, the number of boundary edges of the convex hull, implying that $M + \mathbb{E}[\mathit{nfl}(T)] \geq 3n-3$. We arrived at the main result of this section.

Corollary 5.8. *For any set P of $n \geq 3$ points in general position in the plane it holds that*

$$\mathsf{pg}(P) \geq \sqrt{8}^{\,n-1} \cdot \mathsf{tr}(P). \tag{5.7}$$

Moreover, one may count $\mathsf{pg}(P)$ in time at most $\frac{\mathrm{poly}(n)}{\sqrt{8}^n} \cdot \mathsf{pg}(P)$.

We observe that the bound is tight for $n = 3$. Let us treat the special case of a point set Γ_n in convex position, i.e., when $k = n$. Any triangulation of such a point set has exactly n non-flippable edges, the edges on the boundary of the convex hull of Γ_n, and the remaining $n-3$ edges, the diagonals of the convex n-gon, are flippable. Hence, $\mathbb{E}[nfl(T)] = n = k$ and we cannot improve over the statement of Corollary 5.8 using Theorem 5.7. However, note that $\sqrt{8} \approx 2.828$, and for Γ_n it is known that $\frac{\mathsf{pg}(\Gamma_n)}{\mathsf{tr}(\Gamma_n)} = \Theta^*\left((\frac{3}{2}+\sqrt{2})^n\right)$, where $\frac{3}{2}+\sqrt{2} \approx 2.914$, see Flajolet and Noy [31]. We conjecture that the convex n-gon actually minimizes the fraction $\frac{\mathsf{pg}(P)}{\mathsf{tr}(P)}$ over all sets P of n points in general position.

In the next section we propose a framework for deriving stronger lower bounds on $\mathbb{E}[nfl(T)]$, given that the underlying point set is not in convex position but, conversely, has a triangular convex hull.

5.3 Non-flippable edges in a random graph

In the following assume that P has a triangular convex hull. Actually, the same arguments also work for point sets with at most six points on the convex hull. The basic idea for proving a lower bound on the expected number of non-flippable edges is similar to the method by Sharir and Welzl [72] for estimating the number of degree-3 vertices in a random triangulation. There, every vertex receives an initial charge which it then discharges to vertices of degree 3. Here, however, we want to have each vertex in any triangulation ultimately charge non-flippable edges. If every vertex discharges at least 1 on average and each non-flippable edge receives a charge of at most c, then $\mathbb{E}[nfl(T)]$ is at least the c-th fraction of the total number of vertices.

To make this statement more precise recall that $I^P(P)$ is the set of points in P except for the three extreme points of its convex hull. Then the ground set for our considerations is $I^P(P) \times \mathcal{T}r(P)$ whose elements are called *vints* (*v*ertex-*in*-*t*riangulation). The degree of a vint (p, T) is the degree of the vertex p in the triangulation T. For $i \in \mathbb{N}$, a vint of degree i is called i-vint, and given a fixed triangulation T we denote by $v_i = v_i(T)$ the number of i-vints in $I^P(P) \times \{T\}$. Observe that in any triangulation $v_1 = v_2 = 0$ and $\sum_{i \geq 3} v_i = \left|I^P(P)\right| = n - 3$. Proofs of the following statement can be found in [69, 72].

Lemma 5.9. *Let T be a fixed triangulation and place a charge of $7 - i$ at every i-vint, for $i \geq 3$. Then the weighted sum of charges over all corresponding vints in $I^P(P) \times \{T\}$ is at least $\left|I^P(P)\right|$.*

Proof. Let $d_1, d_2, d_3 \in \mathbb{N}$ denote the degrees of the three extreme vertices in $P \setminus I^P(P)$ of the triangulation T. Then by the Handshaking Lemma, when summing up all degrees in the triangulation T, we get

$$d_1 + d_2 + d_3 + \sum_{i \geq 3} i \cdot v_i = 2(3n - 6) = 6n - 12.$$

We notice that in a triangulation each vertex has degree at least 2, hence $d_1 + d_2 + d_3 \geq 6$ and, therefore, $\sum_{i \geq 3} i \cdot v_i \leq 6n - 18 = 6 \cdot |I^P(P)|$. Thus, the sum of all the vints' charges in the triangulation T is

$$\sum_{i \geq 3} (7 - i) \cdot v_i = 7 \cdot \sum_{i \geq 3} v_i - \sum_{i \geq 3} i \cdot v_i \geq 7 \cdot |I^P(P)| - 6 \cdot |I^P(P)| = |I^P(P)|,$$

i.e., every vint charges at least 1 on average, as desired. □

Note that i-vints with $i \geq 7$ do not receive a positive charge. Hence, it suffices to focus on distributing the charges of 3-, 4-, 5-, and 6-vints. For this we define a relation on the set of vints as in [72]. Let $u = (p_u, T_u)$ and $v = (p_v, T_v)$ be vints then we write $u > v$ if $p_u = p_v$ and there is a flippable edge incident to p_u in T_u such that flipping this edge results in the triangulation T_v. Clearly, in this case u is an $(i+1)$-vint and v an i-vint, for some $i \geq 3$. We denote by $\rightarrow$ the transitive, reflexive closure of $>$. If $u \rightarrow v$ we say that u *may be flipped down to* v.

When discharging we allow every vint to distribute its charge both to lower-degree vints it can be flipped down to and to non-flippable edges. Hereby, a vint (p, T) may only discharge to a non-flippable edge in $E(T)$ that is incident to p or to an edge $qr \in E(T)$ where pqr is a triangle in T. We call such an edge qr a *non-flippable boundary edge of* p. With slight abuse of notation we will also refer to these notions as the (non-flippable) incident and boundary edges of the vint (p, T).

The discharging will be done in such a way that finally there is no positive charge left on any vint. Thus, the sum of charges over all vints in $I^P(P) \times \mathcal{T}r(P)$ which has been distributed among the non-flippable edges is at least $|I^P(P)| \cdot \mathsf{tr}(P)$ by Lemma 5.9. If we can show that during this process a non-flippable edge receives a charge of at most c, then the total number of non-flippable edges in all triangulations of $\mathcal{T}r(P)$ is at least $\frac{1}{c}|I^P(P)| \cdot \mathsf{tr}(P)$. Hence,

$$\mathbb{E}[\mathit{nfl}(T)] \geq \frac{|I^P(P)|}{c}. \tag{5.8}$$

We note that for an edge the property of being non-flippable is by definition equivalent to being incident to a vertex at a, after deleting the edge, reflex angle. In the following we assume the non-flippable edges of the triangulations to be directed towards their endpoint with the reflex angle. Observe that by doing so every non-flippable edge gets directed in a unique way, except for the edges on the boundary of the convex hull.

5.3.1 A simple charging scheme

As an instructive example we will now discuss a simple charging scheme by explicitly stating the distribution of the charges from i-vints to non-flippable edges and other vints. This will result in a first non-trivial lower bound for $\mathbb{E}[nfl(T)]$.

Consider a 3-vint with an initial charge of 4. Since we assumed general position all three incident edges are non-flippable and directed towards the 3-vint, see Figure 5.6(a). Hence, this vint may discharge by equally distributing one third of its charge to the incident non-flippable edges. Note, however, that the 3-vint might still receive charge from higher-degree vints, hence at this point we cannot yet explicitly determine its maximum possible charge to the incoming edges.

(a) 3-vint (b) 4-vint

Figure 5.6: Discharging 3-vints and 4-vints

Observe that a 4-vint always is incident to exactly two non-flippable incoming edges since one of the angles between two non-neighboring edges is always reflex, see Figure 5.6(b). To each of these edges we equally distribute half of the 4-vint's initial charge of 3. In this simple scheme no other vint will discharge to a 4-vint.

When devising the charging scheme for higher-degree vints we further discriminate them according to the number of non-flippable incoming edges. A 5-vint with an initial charge of 2 may occur with two, one or no non-flippable incoming edges, see Figure 5.7. If the 5-vint is incident to at least one non-flippable edge we equally distribute the charge to all such edges.

Note that otherwise none of the five incident edges is directed towards the 5-vint. However, some could still be non-flippable and we cannot directly discharge as we did in the cases before. However, we obtain a 3-vint if we can flip two edges which are not neighbors when considering the edges in clockwise order around the vint. It is always possible to choose two such flippable edges since the boundary edges of the vint define a 5-gon which has at least three convex angles, hence the corresponding edges are flippable. We pass the whole charge of the 5-vint to the resulting 3-vint.

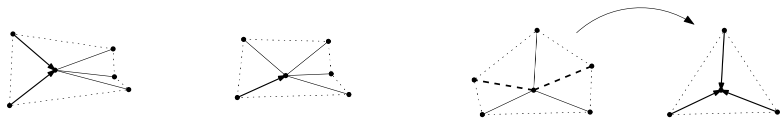

Figure 5.7: Discharging 5-vints

It is crucial to observe that a 3-vint may receive such a charge from at most one 5-vint without non-flippable incoming edges. In order to see this recall that we flipped two non-neighboring edges incident to the 5-vint resulting in two of the three boundary edges of the 3-vint. Assume that two 5-vints flip down to the same 3-vint then one boundary edge is obtained in both flips. Reversing the flip of the common boundary edge results in a 4-vint with two non-flippable incoming edges. If this vint is obtained from a 5-vint without non-flippable edges it is necessary (but not sufficient) to flip the edge of the triangle containing both non-flippable edges of the 4-vint. Hence, there was at most one such 5-vint.

In this scheme a 3-vint will not be charged by any other higher-degree vint, therefore it may end up with a total charge of at most $4 + 2 = 6$. It then discharges uniformly to its three non-flippable incoming edges.

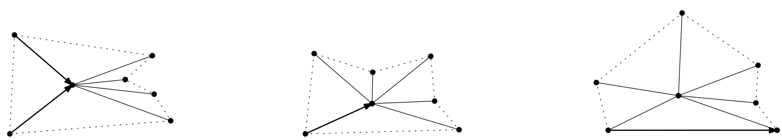

Figure 5.8: Discharging 6-vints with at least one non-flippable edge

Finally, we have to consider the 6-vints with an initial charge of 1. In case such a vint has at least one non-flippable incoming edge we handle it like we did a 5-vint and equally distribute its charge to all those edges. Otherwise we consider the non-flippable boundary edges of the 6-vint, if present, and equally distribute the charge to them, see the right illustration in Figure 5.8.

If, however, a 6-vint neither has non-flippable incoming edges nor non-flippable boundary edges we let it charge to a higher-degree vint for an exception. We flip all boundary edges to obtain a 12-vint of initial charge

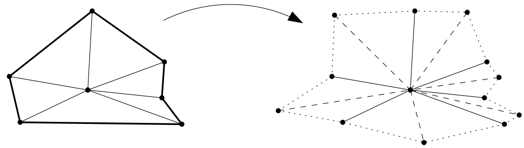

Figure 5.9: Discharging 6-vints without any non-flippable edge

-5 to which we pass the whole charge of the 6-vint, see Figure 5.9. Notice that after all 6-vints have discharged any 12-vint still has a negative charge since it may only be charged by at most two 6-vints. To see this we observe that in order to flip a 12-vint down to such a 6-vint no two neighboring incident edges may be flipped.

Since i-vints with $i \geq 7$ do not have positive charge the average initial charge of 1 by Lemma 5.9 has now been discharged onto the non-flippable edges in the triangulations.

Let us estimate the maximum charge to such a non-flippable edge. For this purpose recall that every non-flippable edge was directed towards a unique vint from which it may be charged, and 6-vints with all incident edges being undirected were the only vints that possibly charge non-flippable boundary edges. Hence, an edge might only receive charge from the one endpoint it is directed to and from at most two 6-vints for which the edge is a non-flippable boundary edge.

In the following we summarize the cases discussed above and list the corresponding maximum charges to an edge depending on the i-vint it is directed towards:

- 3-vint: charge $\leq 1/3 \cdot 6 + 2 = 4$
- 4-vint: charge $\leq 1/2 \cdot 3 + 2 = 3.5$
- 5-vint with two non-flippable edges: charge $\leq 1/2 \cdot 2 + 2 = 3$
- 5-vint with one non-flippable edge: charge $\leq 2 + 2 = 4$
- 6-vint with two non-flippable edges: charge $\leq 1/2 \cdot 1 + 2 = 2.5$
- 6-vint with one non-flippable edge: charge $\leq 1 + 2 = 3$
- i-vint with $i \geq 7$: charge $\leq 2 = 2$.

Therefore, during the discharging of the vints any non-flippable edge receives a charge of at most 4 implying that $\mathbb{E}[nfl(T)] \geq \frac{n-3}{4}$ because of

(5.8). By Theorem 5.7 we have

$$\mathsf{pg}(P) \geq 2^{\frac{3n-6+(n-3)/4}{2}} \cdot \mathsf{tr}(P) = \Omega(2^{13n/8}) \cdot \mathsf{tr}(P) = \Omega(3.084^n) \cdot \mathsf{tr}(P),$$

for any set P of n points in general position with triangular convex hull.

5.3.2 A more elaborate charging scheme

In the following we will improve on the results of the previous subsection. To that end we point out that so far we only allowed 5-vints without non-flippable incoming edges to discharge to a lower-degree 3-vint, but we did not yet take into account that we also could have split a 5-vint's charge among 4-vints it can be flipped down to. Also 6-vints may be flipped down to lower-degree vints and charge them.

Furthermore, recall that a 6-vint with no non-flippable incoming edge may crucially charge its non-flippable boundary edges. In the worst case an edge might receive an additional charge of 2 from such 6-vints. However, for instance no such edge can be directed towards a 3-vint, hence we overestimated the corresponding maximum charge. This shows that there is some potential to improve on the bounds from the discussion above.

In order to generalize the approach for obtaining lower bounds on $\mathbb{E}[\mathit{nfl}(T)]$ note that we actually solved a linear program for determining the way and the amount a vint discharges to non-flippable edges and to other vints, and hence also for the value of the maximum charge.

Indeed, we want to find the smallest value α that is larger than every possible charge to a non-flippable edge, such that there is an initial charge of $7-i$ at every i-vint and after the discharging there is no vint with positive charge left. The corresponding linear program is given below.

$$
\begin{array}{lcl}
\text{minimize} & \alpha & \\
\text{s.t.} & \alpha \geq & \{c_3, c_4, c_{5_2}, c_{5_1}, c_{6_2}, c_{6_1}\} + 2 \cdot b_{6_0}
\end{array}
$$

$$
\begin{array}{rcl@{\qquad}rcl@{\qquad}rcl}
\text{out}3 & \leq & 3 \cdot c_3 & \text{in}3 & \geq & c_{5_0 \to 3} & \text{out}3 & \geq & 4 + \text{in}3 \\
\text{out}4 & \leq & 2 \cdot c_4 & \text{in}12 & \geq & 2 \cdot c_{6_0 \to 12} & \text{out}4 & \geq & 3 + \text{in}4 \\
\text{out}5_2 & \leq & 2 \cdot c_{5_2} & & & & \text{out}5_2 & \geq & 2 + \text{in}5_2 \\
\text{out}5_1 & \leq & c_{5_1} & & & & \text{out}5_1 & \geq & 2 + \text{in}5_1 \\
\text{out}5_0 & \leq & c_{5_0 \to 3} & & & & \text{out}5_0 & \geq & 2 + \text{in}5_0 \\
\text{out}6_2 & \leq & 2 \cdot c_{6_2} & & & & \text{out}6_2 & \geq & 1 + \text{in}6_2 \\
\text{out}6_1 & \leq & c_{6_1} & & & & \text{out}6_1 & \geq & 1 + \text{in}6_1 \\
\text{out}6_0 & \leq & \{b_{6_0}, c_{6_0 \to 12}\} & & & & \text{out}6_0 & \geq & 1 + \text{in}6_0 \\
\text{out}12 & \leq & 0 & & & & \text{out}12 & \geq & -5 + \text{in}12
\end{array}
$$

$$\text{all variables} \geq 0$$

In the following we explain the meaning of variables and constraints. Sets indicated by curly brackets $\{\ldots,\ldots\}$ in the list of constraints are understood as several inequalities of the same form each time replacing one element from the set.

The objective of the linear program is to compute the smallest value for α which is at least the maximum possible charge a non-flippable edge might receive during the discharging of the vints.

The variables c_i (c_{i_j}, resp.) represent the charges to an incoming edge of an i-vint (which has j non-flippable incoming edges), b_{6_0} represents the charge of a 6-vint to a non-flippable boundary edge, and $c_{5_0\to 3}$ ($c_{6_0\to 12}$, resp.) the charge of a 5-vint (6-vint, resp.) with no non-flippable incoming edge to a 3-vint (12-vint, resp.).

We distinguish three types of constraints. First, in the left-most column of constraints for every i-vint (with j non-flippable incoming edges) there is a variable out$i$ (outi_j) that represents the amount of charge that leaves such a vint during the discharging. This amount is upper-bounded by the minimum over all the vint's possibilities to discharge.

Then, in the middle column of constraints there is a variable in$i$ (ini_j) for every i-vint (with j non-flippable incoming edges) that represents the charge received from higher- or lower-degree vints. This additional charge to an i-vint is lower-bounded by the maximum over all possible charges from other vints. Recall that in the charging scheme from the previous subsection there are only charges to 3- and 12-vints, therefore the inequalities for 4-, 5-, and 6-vints are not stated explicitly as all variables are non-negative anyway.

Finally, the right-most column of constraints incorporates the initial charge of $7-i$ at an i-vint and ensures that after discharging there is no positive charge left.

Let us now in a more detailed analysis describe how the vints may distribute their charge in order to obtain a better bound on the maximum charge to an edge. For this purpose first observe that a simple polygon has at least three inner angles which are not reflex. Indeed, since the sum of all angles in an i-gon equals $(i-2)\pi$ there can be at most $i-3$ reflex angles. In particular, an i-vint with j non-flippable incident edges has at least $3-j$ flippable edges. We will use this fact several times.

As it turns out the possibility for a 6-vint to charge a boundary edge does not help for minimizing the charge to non-flippable edges, and thus we also do not need the option to charge a 12-vint. Hence, in the following presentation of the linear program charges will only occur to lower-degree vints or to non-flippable incoming edges.

The linear program we want to solve is of the form

$$\begin{array}{rcl} \text{minimize} & \alpha & \\ \text{s.t.} & \alpha \geq & \{c_3, c_4, c_{5_2}, c_{5_1}, c_{6_2}, c_{6_1}\} \\ & \mathcal{C}_1 \;\; \mathcal{C}_2 \;\; \mathcal{C}_3 & \\ \text{all variables} & \geq & 0, \end{array}$$

where the sets $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$ consist of the three types of constraints we already encountered in the representation of the simple charging scheme as a linear program, and which we now further investigate. For this we will need to introduce some new variables whose relevance will be explained as they appear.

Constraints for discharging vints: $\mathcal{C}_1$

The amount of charge leaving a vint during discharging is captured by the inequalities of $\mathcal{C}_1$. The corresponding constraints for all except the 6-vints without non-flippable incoming edges are given below

$$\begin{array}{rcl} \text{out}3 & \leq & 3 \cdot c_3 \\ \text{out}4 & \leq & 2 \cdot c_4 + c_{4\to 3} \\ \text{out}5_2 & \leq & 2 \cdot c_{5_2} + c_{5_2\to 4} \\ \text{out}5_1 & \leq & c_{5_1} + 2 \cdot c_{5_1\to 4} \\ \text{out}5_0 & \leq & 3 \cdot c_{5_0\to 4} + c_{5_0\to 3} \\ \text{out}6_2 & \leq & 2 \cdot c_{6_2} + c_{6_2\to 5_2} \\ \text{out}6_1 & \leq & c_{6_1} + 2 \cdot c_{6_1\to 5_1} \\ \text{out}6_1 & \leq & c_{6_1} + c_{6_1\to 5_1} + c_{6_1\to 5_2} \\ \text{out}6_1 & \leq & c_{6_1} + 2 \cdot c_{6_1\to 5_2}. \end{array}$$

Let us explain the meaning of these inequalities. Analogously to the simple charging scheme from the previous subsection a 3-vint will equally distribute its charge to its three non-flippable incoming edges.

A 4-vint always has two non-flippable incoming edges and we know that one of the other two edges has to be flippable by the observation we made above. Thus, a 4-vint equally charges its non-flippable incident edges with c_4 and may also charge a lower-degree 3-vint with $c_{4\to 3}$.

By the same reasoning a 5-vint with two non-flippable incoming edges is incident to at least one flippable edge, and hence it is allowed to charge its two non-flippable edges equally with c_{5_2}, and additionally the 4-vint it may be flipped down to with $c_{5_2\to 4}$. A 5-vint with one non-flippable incident edge charges this very edge with c_{5_1}, and by the observation above it may also charge two lower-degree 4-vints with $c_{5_1\to 4}$ each. A 5-vint with no non-flippable incident edge may charge three lower-degree 4-vints equally with $c_{5_0\to 4}$, and since at least two of its flippable incident edges are non-neighboring the vint may also pass $c_{5_0\to 3}$ of its charge to a 3-vint.

When discussing charges from 6-vints to lower-degree vints observe that, clearly, any non-flippable edge incident to the 6-vint will also be non-flippable after flipping down to a lower-degree vint.

Therefore, a 6-vint with two non-flippable incoming edges equally gives charge c_{6_2} to these edges, and by the previous observation it may pass $c_{6_2\to 5_2}$ of its charge to a 5-vint which also has two non-flippable edges. On the other hand it is of course possible that flipping down to a lower-degree vint increases the number of non-flippable edges incident to the new vint. Therefore, a 6-vint with one non-flippable incident edge charges this edge with c_{6_1}, and it may also charge two 5-vints with at least one non-flippable incident edge. The corresponding three inequalities with charges $c_{6_1\to 5_1}$ and $c_{6_1\to 5_2}$ are shown above.

The remaining constraints of $\mathcal{C}_1$ for a 6-vint without non-flippable incoming edges are given by

$$\begin{aligned}
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 3\cdot c_{6_0\to 5_0}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 2\cdot c_{6_0\to 5_0} + c_{6_0\to 5_1}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 2\cdot c_{6_0\to 5_0} + c_{6_0\to 5_2}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + c_{6_0\to 5_0} + 2\cdot c_{6_0\to 5_1}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + c_{6_0\to 5_0} + c_{6_0\to 5_1} + c_{6_0\to 5_2}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + c_{6_0\to 5_0} + 2\cdot c_{6_0\to 5_2}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 3\cdot c_{6_0\to 5_1}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 2\cdot c_{6_0\to 5_1} + c_{6_0\to 5_2}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + c_{6_0\to 5_1} + 2\cdot c_{6_0\to 5_2}\\
\text{out}6_0 &\leq 2\cdot c_{6_0\to 4} + 3\cdot c_{6_0\to 5_2}.
\end{aligned}$$

By the previous observation such a 6-vint is incident to at least three flippable edges, two of which are non-neighboring. In fact, there are two pairs of such non-neighboring edges. In order to see this suppose for contradiction that there are exactly three flippable edges which appear in clockwise order around the 6-vint, call it v for now. Moreover, let v_1, v_2, and v_3 be the vertices adjacent to v which are incident to flippable edges, and v_4, v_5, and v_6 the vertices incident to non-flippable edges. Since the latter edges are non-flippable the interior angles at v_4, v_5, and v_6 of the hexagon $v_1, \ldots, v_6$ are reflex. In turn the sum of the interior angles at v_1, v_2, v_3 is at most π. However, in this case the sum of interior angles at the same vertices but in the quadrilateral $v_1, \ldots, v_3, v$ is also at most π, and hence the angle between v_1v and v_3v is reflex. Thus, the edge in-between v_2v is non-flippable which is a contradiction.

Simultaneously flipping two non-neighboring edges yields a 4-vint which we charge with $c_{6_0\to 4}$, and by the argument above there are at least two such choices. Flipping only one edge of the 6-vint results in a 5-vint with some number of non-flippable incident edges, for which there are at least three possibilities. All such combinations are incorporated in $\mathcal{C}_1$ as shown above.

Charges from higher-degree vints: $\mathcal{C}_2$

For listing the constraints in $\mathcal{C}_2$ which represent the charge received from higher-degree vints we distinguish 3- and 4-vints from the different types of 5-vints in the following. We claim that the former are given by

$$
\begin{aligned}
\text{in3} &\geq c_{5_0\to 3} + 3 \cdot c_{4\to 3} \\
\text{in4} &\geq 4 \cdot c_{6_0\to 4} + c_{5_0\to 4} + 2 \cdot c_{5_1\to 4} + c_{5_2\to 4} \\
\text{in4} &\geq 3 \cdot c_{6_0\to 4} + c_{5_0\to 4} + c_{5_1\to 4} + 2 \cdot c_{5_2\to 4} \\
\text{in4} &\geq 3 \cdot c_{6_0\to 4} + c_{5_0\to 4} + 3 \cdot c_{5_2\to 4} \\
\text{in4} &\geq 2 \cdot c_{6_0\to 4} + 3 \cdot c_{5_1\to 4} + c_{5_2\to 4} \\
\text{in4} &\geq c_{6_0\to 4} + 2 \cdot c_{5_1\to 4} + 2 \cdot c_{5_2\to 4} \\
\text{in4} &\geq c_{5_1\to 4} + 3 \cdot c_{5_2\to 4}.
\end{aligned}
$$

Similarly to the simple charging scheme a 3-vint may only be charged by at most one 5-vint without non-flippable incoming edges. In addition, as seen in the constraint set $\mathcal{C}_1$, it is also possible to receive charge from at most three 4-vints which are obtained by flipping one of the boundary edges.

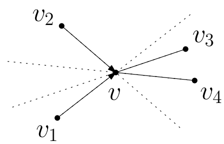

Figure 5.10: A 4-vint v receiving charges from higher-degree vints

As far as the charge to 4-vints is concerned recall that besides 5-vints there is also the possibility for 6-vints without non-flippable edges to charge a lower-degree 4-vint. Note that at most four 5-vints can flip down to a specific 4-vint, and we will distinguish their charges depending on how many 6-vints charge this vint.

Let v be the 4-vint and v_1, v_2, v_3, and v_4 its neighbors such that v_1v and v_2v are the non-flippable incoming edges of v, see Figure 5.10. Observe that the triangulation of a 6-vint which flips down to v cannot simultaneously contain the edges v_2v_3 and v_3v_4, nor the pair v_3v_4 and v_4v_1, since the 6-vint does not have a non-flippable edge. In particular at most four 6-vints may charge v with $c_{6_0\to 4}$.

Notice that if there are exactly four such charges then the different types of 5-vints from which v might also receive charges are determined as follows. If flipping the edge v_1v_2 is possible then v can only become a 5-vint without non-flippable incident edges, flipping v_2v_3 or v_1v_4 may only yield a 5-vint with one non-flippable edge, and finally if v_3v_4 may be flipped a 5-vint with two non-flippable edges is obtained. Therefore, the additional charge to v can be as large as $c_{5_0\to 4} + 2 \cdot c_{5_1\to 4} + c_{5_2\to 4}$.

There are two configurations such that v might be charged by exactly three 6-vints, which also determine the types of the involved 5-vints. For this to happen it is necessary that flipping v_1v_2 yields a 5-vint without non-flippable edges, and flipping v_2v_3 or v_1v_4 at least once results in a 5-vint with two non-flippable edges. Then the additional charge to v from all 5-vints is at most $c_{5_0\to 4} + c_{5_1\to 4} + 2 \cdot c_{5_2\to 4}$ in the first, or $c_{5_0\to 4} + 3 \cdot c_{5_2\to 4}$ in the second case.

If v receives charges from at most two 6-vints then flipping v_1v_2, if possible, has to yield a 5-vint with one non-flippable incident edge. In particular then the number of these 6-vints charging v is exactly the number of edges in $\{v_2v_3, v_1v_4\}$ that yield a 5-vint with one non-flippable edge after flipping. Since flipping v_3v_4 always results in a 5-vint with two non-flippable edges, the additional charge to v can be as large as $3 \cdot c_{5_1\to 4} + c_{5_2\to 4}$ if two 6-vints charge v, or $2 \cdot c_{5_1\to 4} + 2 \cdot c_{5_2\to 4}$ in case of one 6-vint, or $c_{5_1\to 4} + 3 \cdot c_{5_2\to 4}$ if no 6-vint charges v.

The remaining constraints of $\mathcal{C}_2$ for the additional charges to 5-vints are

$$\begin{array}{rcl}
\mathrm{in}5_2 & \geq & c_{6_0\to 5_2} + 4 \cdot c_{6_2\to 5_2} \\
\mathrm{in}5_2 & \geq & c_{6_0\to 5_2} + 3 \cdot c_{6_2\to 5_2} + c_{6_1\to 5_2} \\
\mathrm{in}5_2 & \geq & c_{6_0\to 5_2} + 2 \cdot c_{6_2\to 5_2} + 2 \cdot c_{6_1\to 5_2} \\
\mathrm{in}5_2 & \geq & 4 \cdot c_{6_2\to 5_2} + c_{6_1\to 5_2} \\
\mathrm{in}5_2 & \geq & 3 \cdot c_{6_2\to 5_2} + 2 \cdot c_{6_1\to 5_2} \\
\mathrm{in}5_2 & \geq & 2 \cdot c_{6_2\to 5_2} + 3 \cdot c_{6_1\to 5_2} \\
\mathrm{in}5_1 & \geq & 5 \cdot c_{6_1\to 5_1} \\
\mathrm{in}5_1 & \geq & 4 \cdot c_{6_1\to 5_1} + c_{6_0\to 5_1} \\
\mathrm{in}5_1 & \geq & 3 \cdot c_{6_1\to 5_1} + 2 \cdot c_{6_0\to 5_1} \\
\mathrm{in}5_0 & \geq & 5 \cdot c_{6_0\to 5_0}.
\end{array}$$

Clearly, a particular 5-vint may receive additional charge from at most five distinct 6-vints. Also recall that any non-flippable incident edge to a 6-vint will remain non-flippable when charging a lower-degree vint. However, not all combinations of such charges may actually occur.

Consider a 5-vint with two non-flippable edges that receives charge from 6-vints. By similar arguments as for the charges to 4-vints above it may be observed that at most one of these 6-vints has no non-flippable incoming edge, and not all five but at least two of the 6-vints are incident to two non-flippable edges. All six corresponding inequalities are contained in the constraint set $\mathcal{C}_2$.

As already mentioned a 5-vint with one non-flippable incident edge cannot receive charge from a 6-vint with two incident non-flippable edges. Out of the five 6-vints charging such a vint at most two have no non-flippable incident edge. This allows for three possibilities to receive charge from the higher-degree vints.

Finally, the last inequality shows that a 5-vint without non-flippable incident edges may only receive charges from corresponding 6-vints.

Initial charges and consumption: $\mathcal{C}_3$

The constraints of $\mathcal{C}_3$ introduce the initial charges and make sure that after discharging there is no positive charge left on any vint

$$\begin{aligned}
\mathrm{out3} &\geq 4 + \mathrm{in3} \\
\mathrm{out4} &\geq 3 + \mathrm{in4} \\
\mathrm{out5}_2 &\geq 2 + \mathrm{in5}_2 \\
\mathrm{out5}_1 &\geq 2 + \mathrm{in5}_1 \\
\mathrm{out5}_0 &\geq 2 + \mathrm{in5}_0 \\
\mathrm{out6}_2 &\geq 1 + \mathrm{in6}_2 \\
\mathrm{out6}_1 &\geq 1 + \mathrm{in6}_1 \\
\mathrm{out6}_0 &\geq 1 + \mathrm{in6}_0.
\end{aligned}$$

Now, we are ready to compute the optimum of the linear program we derived. Recall that we want to find the minimum α which is at least as big as $\max\{c_3, c_4, c_{5_2}, c_{5_1}, c_{6_2}, c_{6_1}\}$ subject to the constraint sets $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$, and all variables are non-negative. Solving this linear program yields $\alpha = \frac{37}{18}$ as an optimal value for the maximum possible charge to a non-flippable edge. Together with (5.8) and Theorem 5.7 we obtain the following result.

Theorem 5.10. *For any set P of $n \geq 3$ points in the plane in general position with triangular convex hull,* $\mathbb{E}[\text{nfl}(T)] \geq \frac{18(n-3)}{37}$, *and hence*

$$\mathsf{pg}(P) \geq \Omega(2^{129n/74}) \cdot \mathsf{tr}(P) = \Omega(3.347^n) \cdot \mathsf{tr}(P).$$

The lower bound $\Omega(3.347^n)$ on the fraction $\frac{\mathsf{pg}(P)}{\mathsf{tr}(P)}$ from Theorem 5.10 holds for every point set P in general position. This compares to a value of $O(4.855^n)$ which is obtained by a particular point set, the previously mentioned double zig-zag chain D_n due to Aichholzer et al. [8] who calculate $\mathsf{pg}(D_n) = \Theta^*(41.189^n)$ and $\mathsf{tr}(D_n) = \Theta^*(8.485^n)$. With (5.6) this point set also implies the bound $\mathbb{E}[\mathit{nfl}(T)] \leq 1.559n$.

Actually, the convex hull of D_n is a quadrangle but the point set may be slightly modified in order to have triangular convex hull while keeping the asymptotic behavior of the number of crossing-free graphs and triangulations.

To see this assume very flat zig-zag chains and rotate the upper chain clockwise and the lower chain counter-clockwise without changing the point configuration. Then place an additional point very far to the right extending both chains in a convex way, i.e., points on the convex hull of a single chain remain on the hull when adding the special point.

We conclude by mentioning that in the corresponding Figure 5.11 for the sake of a reasonable drawing this additional point is in fact too close to the construction.

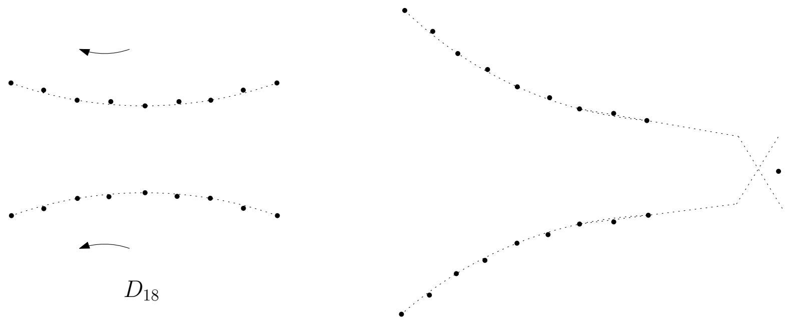

Figure 5.11: The double zig-zag chain D_{18} and its slightly modified version

I don't see enough DoTs - more DoTs now!
Throw more DoTs, more DoTs, more DoTs!
Come on more DoTs!!

Dives, from Wipe Club

6

Slight Perturbation causes Exponential Change

We show that there is a constant $\beta \geq \frac{1}{144} > 0$ such that, for any set P of $n \geq 5$ points in general position in the plane, a crossing-free geometric graph on P that is chosen uniformly at random contains, in expectation, at least $(\frac{1}{2} + \beta)M$ edges, where M denotes the number of edges in any triangulation of P.

From this we derive the first non-trivial upper bound on the number of crossing-free geometric graphs on P of the form $c^n \cdot \mathsf{tr}(P)$; That is, at most a fixed exponential in n times the number of triangulations of P. The trivial upper bound of $2^M \cdot \mathsf{tr}(P)$, or $c = 2^{M/n}$, follows by taking all subsets of edges of each triangulation. If the convex hull of P is triangular, then $M = 3n - 6$, and we obtain $c < 7.980$.

Upper bounds for the number of crossing-free geometric graphs on planar point sets have so far applied the trivial 8^n factor to the bound for triangulations; We slightly decrease the resulting bound to $O(343.106^n)$.

This is joint work with Jack Snoeyink and Emo Welzl [64].

6.1 Introduction

Let P be a finite set of points in the plane. We define crossing-free geometric graphs and triangulations on P as in the previous chapter, and write n for the number of points in P and k for the number of points on the boundary of its convex hull. For the sake of a clean presentation of the results we will again assume that $n \geq k \geq 3$. Recall from identity (5.1) that any triangulation of P contains exactly $M := 3n - k - 3$ edges. In contrast to the previous chapter we say that P is in *general position* if no three points in P are collinear.

At the end of our discussion we will provide an upper bound for $\mathsf{pg}(P)$ the total number of plane graphs on P. This quantity never exceeds a fixed exponential in n which is a result first established in 1982 by Ajtai et al. [12] with 10^{13} as base of the exponential. Further progress [75, 70, 26, 69] in this area of research was mainly motivated by deriving better bounds for $\mathsf{tr}(P)$, the number of triangulations on P, where the currently best known upper bound for $\mathsf{tr}(P)$ stands at 43^n due to Sharir and Welzl [72]. We will arrive at an upper bound of 343.106^n for $\mathsf{pg}(P)$.

While, clearly, upper bounds on the total number of crossing-free geometric graphs on a point set P also apply to specific classes of plane graphs (e.g. spanning connected graphs, polygonizations, perfect matchings, and spanning trees, to name just a few), better bounds for these classes are known. For a recent and detailed list of such results we refer to [8, 71].

We will estimate $\mathsf{pg}(P)$ in terms of $\mathsf{tr}(P)$. Since every crossing-free geometric graph is contained in some triangulation, and every triangulation has 2^M subgraphs, we readily have

$$\mathsf{pg}(P) \leq 2^M \cdot \mathsf{tr}(P) \leq 8^n \cdot \mathsf{tr}(P). \tag{6.1}$$

The upper bound is tight in the following example: Consider a point set with triangular convex hull such that all interior points lie on a common line containing one of the three extreme points, see Figure 6.1. Then there is a unique triangulation with exactly $2^{3n-6} = \Theta(8^n)$ crossing-free subgraphs; These subgraphs constitute the set of all plane graphs on the point set.

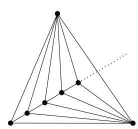

Figure 6.1: Points not in general position with a unique triangulation

It is therefore surprising that a small perturbation of the points to general position causes the ratio between the number of crossing-free graphs and that of triangulations to drop exponentially. We show that for any set P of at least five points in general position $\mathsf{pg}(P) \leq 2^{\gamma \cdot M} \cdot \mathsf{tr}(P)$ holds, with $\gamma < 1$.

In order to derive this bound we argue in Section 6.2 that the expected number of edges in a crossing-free geometric graph on P chosen uniformly at random can be significantly bounded away from $\frac{M}{2}$. In Section 6.3 we prove via a mean vs. median argument that crossing-free geometric graphs with many edges account for a large fraction of all crossing-free graphs. This way, for points in general position we are able to improve on the pessimistic 2^M factor from estimate (6.1) by a factor exponential in n.

6.2 Edges in a random crossing-free graph

The goal of this section is to estimate the expected number of edges in a crossing-free geometric graph drawn uniformly at random from the set of all crossing-free graphs on a given point set P. We will refer by $e(G)$ to the number of edges in a fixed graph G as well as to the random variable when choosing a crossing-free graph G uniformly at random.

We define a directed graph $\mathcal{D} = \mathcal{D}(P)$ on the set of all crossing-free geometric graphs on P and for two nodes $G, H \in V(\mathcal{D})$ introduce a directed arc from G to H if and only if $E(G) \subseteq E(H)$ and $e(G) = e(H) - 1$. For instance, the empty graph has $\binom{n}{2}$ outgoing but zero incoming arcs, while any triangulation of P has no outgoing but M incoming arcs. Expectations of random variables are in the following understood with respect to the uniform distribution over the set of all crossing-free geometric graphs.

Proposition 6.1. *For any point set in general position in the plane*

$$\frac{M}{2} \leq \mathbb{E}[e(G)] \leq M.$$

Proof. Clearly, $e(G) \leq M$ holds for any graph G which immediately implies the upper bound. For proving the lower bound denote by $\deg^-(G)$ and $\deg^+(G)$ the in- and out-degree of G in the directed graph $\mathcal{D}$. Notice that

$$\mathbb{E}[e(G)] = \mathbb{E}\left[\deg^-(G)\right] = \mathbb{E}\left[\deg^+(G)\right],$$

where the first identity holds since $e(G) = \deg^-(G)$ for any graph G, and the second equality is true since by definition both sides represent the number of arcs divided by the number of vertices of $\mathcal{D}$. Then using $\deg(G) := \deg^-(G) + \deg^+(G)$ and by linearity of expectation we obtain

$$2 \cdot \mathbb{E}[e(G)] = \mathbb{E}\left[\deg^-(G)\right] + \mathbb{E}\left[\deg^+(G)\right] = \mathbb{E}[\deg(G)]\,. \tag{6.2}$$

Consider some triangulation T with M edges that contains G as a subgraph. Then any edge $e \in E(T)$ either corresponds to an incoming arc of G if $e \in E(G)$, or to an outgoing arc of G if $e \notin E(G)$. Thus, $\deg(G) \geq M$, which completes the proof of the first inequality. □

Note that the lower bound is tight only for $n = 3$ points, or $n = 4$ points with a triangular convex hull. In all other cases the underlying point set allows for more than one triangulation. Hence $\binom{n}{2}$, the degree of the empty graph in $\mathcal{D}$, is strictly larger than M which implies $\mathbb{E}[\deg(G)] > M$ and in turn $\mathbb{E}[e(G)] > \frac{M}{2}$.

Consider a set of points with $n = k = 4$; Figure 6.2 shows the vertices of the corresponding directed graph $\mathcal{D}$ and the edges incident to one graph consisting of a single diagonal. As is easily verified, for a crossing-free graph G chosen uniformly at random $\mathbb{E}[e(G)] = \frac{8}{3} > \frac{5}{2}$.

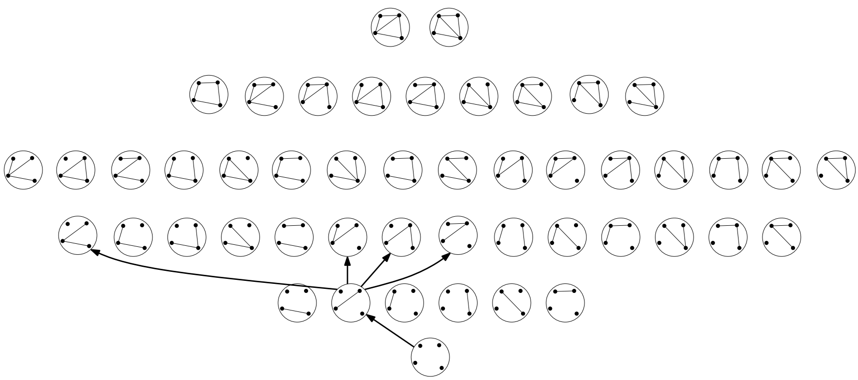

Figure 6.2: The directed graph $\mathcal{D}$ on four points

From now on we focus on point sets containing at least five points and show that there is a constant $\beta > 0$ such that $\mathbb{E}[e(G)] \geq (\frac{1}{2} + \beta)M$.

Let G be a crossing-free geometric graph on P and let $e \notin E(G)$ be an edge corresponding to an outgoing arc of G in $\mathcal{D}$, hence adding e to G again yields a crossing-free graph. If every triangulation containing G also contains e, then we call e *forced* for G. Otherwise, we call edge e *optional* for G. Edges from the convex hull missing in G are always forced for G. For instance, if $n = 4$ and $k = 3$ then every edge not in $E(G)$ is forced.

Arcs in $\mathcal{D}$ are labeled forced or optional to match their corresponding edge. In the previous example of Figure 6.2, the drawn incoming arc is optional whereas the four outgoing arcs are forced. For a graph G we may define its *forced degree*, $\mathrm{fdeg}(G)$, and its *optional degree*, $\mathrm{odeg}(G)$. We write $\mathrm{fdeg}^+(G)$ and $\mathrm{odeg}^+(G)$ for the corresponding forced and optional out-degree, respectively.

Lemma 6.2. *Given a crossing-free geometric graph G, adding the set of all its forced edges results in a crossing-free geometric graph $\overline{G}$ without forced edges.*

Proof. Let $E_f(G)$ be the set of forced edges for G. Then $\overline{G}$ is crossing-free since $E(\overline{G}) = E(G) \cup E_f(G)$ is the set of edges which are present in all triangulations containing G. In other words, $\overline{G}$ is the largest graph that is a common subgraph of all triangulations that contain G. Now assume that there is a forced edge e for $\overline{G}$. Since e was not forced for G (otherwise $e \in E_f(G) \subseteq E(\overline{G})$), there is a triangulation containing G but not e. This is a contradiction since $\overline{G}$ is a subgraph of this triangulation and e must be an edge of every triangulation containing $\overline{G}$. □

For $\overline{G}$ as obtained in Lemma 6.2 let $u(G) := M - e(\overline{G})$ be the number of edges we need to add to $\overline{G}$ in order to obtain a triangulation. Then by definition

$$M - u(G) = e(\overline{G}) = e(G) + \mathrm{fdeg}^+(G) = \deg^-(G) + \mathrm{fdeg}^+(G).$$

Having equation (6.2) in mind we are interested in the expected value of $\deg(G) - M =: \mathrm{excess}(G)$ which, given the identity above, we can rewrite as

$$\mathrm{excess}(G) = \deg^+(G) + \deg^-(G) - M = \mathrm{odeg}^+(G) - u(G), \qquad (6.3)$$

using $\deg^+(G) = \mathrm{fdeg}^+(G) + \mathrm{odeg}^+(G)$. We want to show that $\mathrm{odeg}^+(G)$ is large compared to $u(G) = u(\overline{G})$. Assume $u(\overline{G}) > 0$ and let e be an edge corresponding to an outgoing arc of $\overline{G}$ in $\mathcal{D}$ then any triangulation containing $\overline{G}$ but not e must necessarily contain an edge f that crosses e, otherwise we could add e to the triangulation, contradicting its maximality. Thus, also f corresponds to an outgoing arc of $\overline{G}$ in $\mathcal{D}$ and clearly both arcs, e and f, are optional for G. This suggests that one can find at least $2u(G)$ optional outgoing arcs of G.

However, establishing this takes some care since two edges e_1 and e_2, corresponding to outgoing arcs of $\overline{G}$, could be crossed by the same edge f in a triangulation containing $\overline{G}$. This would require us to find a compensating fourth optional outgoing arc of G. We repeatedly apply the following result by Aichholzer et al. [5] to overcome this issue.

Lemma 6.3 (Aichholzer et al. [5]). *Let P be a finite set of points in the plane and consider two triangulations T and T' of P. There exists a perfect matching between the edges of T and T', with the property that matched edges either cross or are identical.*

Lemma 6.4. *For any crossing-free graph G we have* $\operatorname{odeg}^+(G) \geq 2u(G)$.

Proof. In order to prove this claim it suffices to construct a matching C on the set of optional outgoing arcs of G in $\mathcal{D}$ such that $|C| = u(G)$. Let $\overline{G}$ be the graph obtained from G as in Lemma 6.2, T_1 a triangulation containing $\overline{G}$ and define $E_1 := E(T_1) \setminus E(\overline{G})$. We now match edges in E_1, a set of $u(G)$ optional edges, with other edges optional for G.

Initially $C = \emptyset$. For $i \geq 1$, assume that $E_i := E(T_1) \cap \ldots \cap E(T_i) \setminus E(\overline{G})$ is not empty, where T_j, for $1 \leq j \leq i$, are triangulations constructed so far. Let $e \in E_i$. Since e is not forced for $\overline{G}$ there is an edge f crossing e which is also not forced for $\overline{G}$. Let T_{i+1} be a triangulation containing $\overline{G}$ and f. Lemma 6.3 gives a perfect matching between $E(T_1)$ and $E(T_{i+1})$. Consider the matching partners of $E_i \setminus E(T_{i+1}) \subset E(T_1)$. These are edges of $E(T_{i+1}) \setminus \bigcup_{j=1}^{i} E(T_j)$, since they cross edges of $E_i \subset \bigcap_{j=1}^{i} E(T_j)$. We observe that the matched crossing pairs correspond to optional outgoing arcs of G. By construction these pairs are disjoint from all pairs in C, hence we may safely add them to C. In particular, all edges in E_i that cross f will be added to C in some matching pair. Now, we continue with E_{i+1} unless it is empty.

Note that by construction $e \in E_i \setminus E(T_{i+1})$ which implies $|E_{i+1}| < |E_i|$. Furthermore, $|E_i| - |E_{i+1}|$ is exactly the number by which $|C|$ increases in the i-th round. Thus, the process terminates eventually with a matching of size $|E_1| = u(G)$. □

Substitute this estimate $u(G) \leq \frac{1}{2} \cdot \operatorname{odeg}^+(G)$ into identity (6.3) to obtain $\operatorname{excess}(G) \geq \frac{\operatorname{odeg}^+(G)}{2}$. Similarly to the proof of Proposition 6.1 we find $2 \cdot \mathbb{E}\left[\operatorname{odeg}^+(G)\right] = \mathbb{E}\left[\operatorname{odeg}(G)\right]$ which yields

$$\mathbb{E}\left[\operatorname{excess}(G)\right] \geq \frac{1}{4} \cdot \mathbb{E}\left[\operatorname{odeg}(G)\right].$$

If G happens to be a triangulation then by definition the optional edges for G are exactly the flippable edges in G.

Theorem 6.5 (Hurtado et al. [40]). *Any triangulation of a collection of n points on the plane contains at least $\frac{n-4}{2}$ flippable edges. The bound is tight.*

This is the kind of result we are heading for, however, we need a corresponding statement for every graph on the underlying point set not for triangulations only.

Lemma 6.6. *For any crossing-free graph G it holds that*

$$\mathrm{odeg}(G) \geq \frac{n-4}{2} + u(G).$$

Proof. Extend the graph G to obtain a triangulation T with $E(G) \subseteq E(T)$ and let $e \in E(T)$ be a flippable edge, i.e., e corresponds to an optional incoming arc of T in $\mathcal{D}$. By Theorem 6.5 there are at least $\frac{n-4}{2}$ such edges. If $e \in E(G)$ then e adds to $\mathrm{odeg}(G)$ as (optional) incoming arc of G in $\mathcal{D}$. If $e \notin E(G)$ then e adds to $\mathrm{odeg}(G)$ as (optional) outgoing arc of G.

Moreover, we know by Lemma 6.4 that G has at least $2u(G)$ optional outgoing edges. Since at most $u(G)$ of them are contained in T the remaining edges clearly add to $\mathrm{odeg}(G)$. □

We found that

$$\mathbb{E}\left[\mathrm{excess}(G)\right] \geq \frac{n-4}{8} + \frac{\mathbb{E}\left[u(G)\right]}{4},$$

unfortunately in general we were not able to give a lower bound for $\mathbb{E}[u(G)]$ other than the trivial one, $\mathbb{E}[u(G)] \geq 0$. Since $\deg(G) = M + \mathrm{excess}(G)$ we have $\mathbb{E}[e(G)] = \frac{1}{2} \cdot (M + \mathbb{E}[\mathrm{excess}(G)])$ due to equation (6.2).

Theorem 6.7. *Let P be a set of $n \geq 4$ points in general position in the plane with k points on the boundary of the convex hull and $M = 3n - k - 3$ edges in a triangulation. Then*

$$\mathbb{E}\left[e(G)\right] \geq \frac{M}{2} + \frac{n-4}{16} = \frac{25n - 8k - 28}{16}. \tag{6.4}$$

While the statement of the theorem is also true for $n = 3$ it is only of moderate interest because of Proposition 6.1. For $n \geq 5$, however, choosing $\beta = \frac{1}{144}$ the bound from the theorem above shows that $\mathbb{E}[e(G)] \geq (\frac{1}{2} + \beta)M$ for a crossing-free geometric graph drawn uniformly at random.

6.3 The number of plane graphs

In this section we provide an argument showing how the lower bound from (6.4) yields an upper bound on the total number of crossing-free geometric graphs a set of n points can have. For this purpose we will need two classic estimates that frequently appear in combinatorics. We also make use of the binary entropy function

$$\mathcal{H}(t) := -t\log_2 t - (1-t)\log_2(1-t), \quad t \in \,]0,1[,$$

with continuous extension $\mathcal{H}(0) = \mathcal{H}(1) = 0$. The function is symmetric around $\frac{1}{2}$ where it attains its maximum $\mathcal{H}(\frac{1}{2}) = 1$, and it strictly increases on the interval $]0, \frac{1}{2}]$, and thus strictly decreases on $[\frac{1}{2}, 1[$.

The following estimate for sums of binomial coefficients is standard, we include a proof for the sake of the thesis being self-contained.

Lemma 6.8. *For any $\ell \in \mathbb{N}$ and $t \in \; [0, \frac{1}{2}]$*

$$\sum_{i=0}^{\lfloor t\cdot\ell\rfloor} \binom{\ell}{i} \leq 2^{-\mathcal{H}(t)\cdot\ell}.$$

Proof. Since $t \in \; [0, \frac{1}{2}]$ it follows that $\frac{t}{1-t} \leq 1$ and hence

$$\begin{aligned}\sum_{i=0}^{\lfloor t\cdot\ell\rfloor} \binom{\ell}{i} \cdot 2^{-\mathcal{H}(t)\cdot\ell} &= \sum_{i=0}^{\lfloor t\cdot\ell\rfloor} \binom{\ell}{i} \cdot t^{t\ell}(1-t)^{(1-t)\ell} \\ &= \sum_{i=0}^{\lfloor t\cdot\ell\rfloor} \binom{\ell}{i} \cdot (1-t)^{\ell}\left(\frac{t}{1-t}\right)^{t\ell} \\ &\leq \sum_{i=0}^{\lfloor t\cdot\ell\rfloor} \binom{\ell}{i} \cdot (1-t)^{\ell}\left(\frac{t}{1-t}\right)^{i}.\end{aligned}$$

Now, this last term is surely upper-bounded by

$$\sum_{i=0}^{\ell} \binom{\ell}{i} \cdot t^i \cdot (1-t)^{\ell-i} = (t + (1-t))^{\ell} = 1,$$

using the Binomial theorem, which concludes the proof of the lemma. □

Another tool we will employ in the upcoming derivation of our main result is the following inequality.

Lemma 6.9 (Markov's Inequality)**.** *Let X be a non-negative random variable of expected value $\mathbb{E}[X]$ then for any real $\lambda > 0$*

$$\mathbb{P}[X \geq \lambda] \leq \frac{\mathbb{E}[X]}{\lambda}.$$

Proof. By definition of the indicator function, $X \geq \lambda \cdot \mathbb{1}_{[X \geq \lambda]}$, and since the expected value is monotone we obtain

$$\mathbb{E}[X] \geq \mathbb{E}\left[\lambda \cdot \mathbb{1}_{[X \geq \lambda]}\right] = \lambda \cdot \mathbb{P}[X \geq \lambda]. \qquad \square$$

We are now able to conclude with the upper bound for the total number $\mathsf{pg}(P)$ of crossing-free geometric graphs.

Theorem 6.10. *Let P be a set of $n \geq 4$ points in general position in the plane with k points on the boundary of the convex hull. Moreover, define $M = 3n - k - 3$ and $\mu = \mathbb{E}[e(G)]$. Then*

$$\mathsf{pg}(P) \leq M \cdot 2^{\mathcal{H}(\frac{\lfloor \mu \rfloor}{M})M} \cdot \mathsf{tr}(P).$$

Proof. We will first show that the crossing-free geometric graphs with at least $\lfloor \mu \rfloor$ edges, which we denote by $\mathsf{pg}_{\geq \lfloor \mu \rfloor}(P)$, form a large fraction of all crossing-free graphs. In order to achieve this we provide a lower bound for

$$\frac{\mathsf{pg}_{\geq \lfloor \mu \rfloor}(P)}{\mathsf{pg}(P)} = \mathbb{P}[e(G) \geq \lfloor \mu \rfloor].$$

For this purpose we note that the random variable $e(G)$ only takes integer values, hence

$$\begin{aligned} \mathbb{P}[e(G) \geq \lfloor \mu \rfloor] &= \mathbb{P}[e(G) > \mu - 1] \\ &= 1 - \mathbb{P}[e(G) \leq \mu - 1] \\ &= 1 - \mathbb{P}[M - e(G) \geq M - (\mu - 1)]. \end{aligned}$$

Clearly, the random variable $M - e(G)$ is non-negative. Therefore, we may estimate the last probability with Markov's inequality from Lemma 6.9 to obtain

$$\begin{aligned} \mathbb{P}[e(G) \geq \lfloor \mu \rfloor] &\geq 1 - \frac{\mathbb{E}[M - e(G)]}{M - (\mu - 1)} \\ &= 1 - \frac{M - \mu}{M - (\mu - 1)} \\ &= \frac{1}{M - (\mu - 1)} \geq \frac{1}{M}, \end{aligned}$$

since $\mu \geq 1$. Thus, we have $\mathsf{pg}_{\geq \lfloor \mu \rfloor}(P) \geq \frac{\mathsf{pg}(P)}{M}$.

On the other hand, let us for the moment fix a triangulation and count the number of its crossing-free subgraphs with at least $\lfloor \mu \rfloor$ edges

$$\sum_{m=\lfloor \mu \rfloor}^{M} \binom{M}{m} = \sum_{m'=0}^{M-\lfloor \mu \rfloor} \binom{M}{m'} = \sum_{m'=0}^{M\left(1-\frac{\lfloor \mu \rfloor}{M}\right)} \binom{M}{m'}.$$

In order to employ Lemma 6.8 for estimating this sum we would need that $1 - \frac{\lfloor \mu \rfloor}{M}$ is in $[0, \frac{1}{2}]$. Recall Proposition 6.1 stating that $\frac{M}{2} \leq \mu \leq M$. Clearly, for even M also $\frac{M}{2} \leq \lfloor \mu \rfloor \leq M$ holds, and thus $0 \leq 1 - \frac{\lfloor \mu \rfloor}{M} \leq \frac{1}{2}$. If M is odd and $n \geq 12$ then by Theorem 6.7 we have

$$\mu \geq \frac{M}{2} + \frac{n-4}{16} \geq \frac{M+1}{2} \in \mathbb{N},$$

implying $\lfloor \mu \rfloor \geq \frac{M+1}{2}$ which also satisfies the condition. For the remaining values of M and n, that is when $4 \leq n \leq 11$ and thus $5 \leq M \leq 27$ and M odd, the factor might be larger than $\frac{1}{2}$. Despite this fact one verifies by numerical calculation that the estimate from Lemma 6.8 still holds. Therefore, we conclude

$$\sum_{m'=0}^{M\left(1-\frac{\lfloor \mu \rfloor}{M}\right)} \binom{M}{m'} \leq 2^{\mathcal{H}\left(1-\frac{\lfloor \mu \rfloor}{M}\right)M} = 2^{\mathcal{H}\left(\frac{\lfloor \mu \rfloor}{M}\right)M}.$$

Since every crossing-free graph is contained in some triangulation, we may estimate the total number of crossing-free graphs with at least $\lfloor \mu \rfloor$ edges by summing over all triangulations and obtain

$$\mathsf{pg}_{\geq \lfloor \mu \rfloor}(P) \leq 2^{\mathcal{H}\left(\frac{\lfloor \mu \rfloor}{M}\right)M} \cdot \mathsf{tr}(P).$$

Lastly, lower and upper bound on $\mathsf{pg}_{\geq \lfloor \mu \rfloor}(P)$ imply the statement of the theorem. □

In fact the estimate of Theorem 6.10 also holds for $n = 3$ which may be easily verified. However, the proof above does not go through in this case since the bound from Lemma 6.8 cannot be applied when $n = k = M = 3$.

Recall that the binary entropy function is strictly decreasing on the interval $[\frac{1}{2}, 1]$, thus the lower bound $\mu \geq \frac{M}{2} + \frac{n-4}{16}$ from equation (6.4) comes in quite handy. A numerical analysis shows that for $n \geq 3596$ the bound from Theorem 6.10 beats the trivial estimate (6.1).

It is not hard to see that a triangular convex hull of the point set P minimizes the lower bound on $\frac{\lfloor\mu\rfloor}{M}$ in (6.4) and therefore maximizes $\mathcal{H}\left(\frac{\lfloor\mu\rfloor}{M}\right)$, for n sufficiently large. Thus, we find that for any set P

$$\mathsf{pg}(P) = O\left(n\,2^{\mathcal{H}\left(\frac{25}{48}\right)3n}\right)\cdot \mathsf{tr}(P) = O(7.9792^n)\cdot \mathsf{tr}(P).$$

We mentioned that there are at most 43^n triangulations on a set of n points [72]. Hence, n points in the plane allow for at most $O(343.106^n)$ crossing-free geometric graphs.

Corollary 6.11. *For a set P of n points in general position in the plane,*

$$\mathsf{pg}(P) = O(7.980^n)\cdot \mathsf{tr}(P) \quad \textit{and} \quad \mathsf{pg}(P) = O(343.106^n).$$

Recall from Chapter 1 the definition of the single chain S_n with triangular convex hull, and notice that every edge incident to the tip vertex must be present in any triangulation of S_n. The remaining $n-1$ points are in convex position which enables us to exactly compute $\mathsf{tr}(S_n) = \mathsf{tr}(\Gamma_{n-1})$ and $\mathsf{pg}(S_n) = 2^{n-1}\mathsf{pg}(\Gamma_{n-1})$. The corresponding asymptotics for Γ_n are known due to [31] and we find $\frac{\mathsf{pg}(S_n)}{\mathsf{tr}(S_n)} = \Theta^*\left((3+2\sqrt{2})^n\right) = \Theta^*(5.828^n)$ as lower bound on the fraction $\frac{\mathsf{pg}(P)}{\mathsf{tr}(P)}$ any set on n points can achieve.

For the other extreme case of Γ_n, a set of n points in convex position, we have $k = n$ and $M = 2n-3$, thus Theorem 6.10 gives

$$\frac{\mathsf{pg}(\Gamma_n)}{\mathsf{tr}(\Gamma_n)} = O\left(n\,2^{\mathcal{H}\left(\frac{17}{32}\right)2n}\right) = O(3.985^n).$$

This bound can slightly be improved to $O\big(n\,2^{\mathcal{H}\left(\frac{9}{16}\right)2n}\big) = O(3.938^n)$ using our previously developed machinery. We simply notice that any triangulation of the convex n-gon has exactly $n-3$ diagonals, i.e., flippable edges. The corresponding estimate in Lemma 6.6 becomes $\mathrm{odeg}(G) \geq n-3$ and in turn the following statements yield $\mu \geq \frac{M}{2} + \frac{n-3}{8}$.

Actually, for points in convex position the expected number of edges $\mu = \mathbb{E}[e(G)]$ is known to be

$$\mathbb{E}[e(G)] = \frac{1+\sqrt{2}}{2}n\cdot\big(1+o(1)\big),$$

due to Flajolet and Noy [31] and Bernasconi et al. [19] who also provide tight estimates for the tail probabilities. Theorem 6.10 immediately yields

$$\frac{\mathsf{pg}(\Gamma_n)}{\mathsf{tr}(\Gamma_n)} = O\left(n\,2^{\mathcal{H}\left(\frac{1+\sqrt{2}}{4}(1+o(1))\right)2n}\right) = O(3.831^n).$$

These bounds on the fraction $\frac{\mathsf{pg}(\Gamma_n)}{\mathsf{tr}(\Gamma_n)}$ compare to the exact asymptotic behavior $\Theta^*\left((\frac{3}{2}+\sqrt{2})^n\right)$, where $\frac{3}{2}+\sqrt{2} \approx 2.914$.

Bibliography

[1] Bernardo Manuel Ábrego and Silvia Fernández-Merchant. Geometric drawings of K_n with few crossings. *J. Combin. Theory Ser. A*, 114(2):373–379, 2007.

[2] Oswin Aichholzer, Franz Aurenhammer, Clemens Huemer, and Hannes Krasser. Transforming spanning trees and pseudo-triangulations. *Inform. Process. Lett.*, 97(1):19–22, 2006.

[3] Oswin Aichholzer, Franz Aurenhammer, and Ferran Hurtado. Sequences of spanning trees and a fixed tree theorem. *Comput. Geom. Theory Appl.*, 21(1):3–20, 2002.

[4] Oswin Aichholzer, Franz Aurenhammer, and Hannes Krasser. Enumerating order types for small point sets with applications. In *SCG '01: Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, pages 11–18, New York, NY, USA, 2001. ACM.

[5] Oswin Aichholzer, Franz Aurenhammer, Günter Rote, and Michael Taschwer. Triangulations intersect nicely. In *SCG '95: Proceedings of the Eleventh Annual Symposium on Computational Geometry*, pages 220–229, New York, NY, USA, 1995. ACM.

[6] Oswin Aichholzer, Sergey Bereg, Adrian Dumitrescu, Alfredo García, Clemens Huemer, Ferran Hurtado, Mikio Kano, Alberto Márquez, David Rappaport, Shakhar Smorodinsky, Diane Souvaine, Jorge Urrutia, and David Ronald Wood. Compatible geometric matchings. *Comput. Geom. Theory Appl.*, 42(6-7):617–626, 2009.

[7] Oswin Aichholzer, Jesús García, David Orden, and Pedro Ramos. New lower bounds for the number of ($\leq k$)-edges and the rectilinear crossing number of K_n. *Discrete Comput. Geom.*, 38(1):1–14, 2007.

[8] Oswin Aichholzer, Thomas Hackl, Clemens Huemer, Ferran Hurtado, Hannes Krasser, and Birgit Vogtenhuber. On the number of plane

graphs. In *SODA '06: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 504–513, New York, NY, USA, 2006. ACM.

[9] Oswin Aichholzer, Ferran Hurtado, and Marc Noy. On the number of triangulations every planar point set must have. In *CCCG '01: Proceedings of the 13th Annual Canadian Conference on Computational Geometry*, pages 13–16, 2001.

[10] Oswin Aichholzer, Ferran Hurtado, and Marc Noy. A lower bound on the number of triangulations of planar point sets. *Comput. Geom. Theory Appl.*, 29(2):135–145, 2004.

[11] Oswin Aichholzer and Klaus Reinhardt. A quadratic distance bound on sliding between crossing-free spanning trees. *Comput. Geom. Theory Appl.*, 37(3):155–161, 2007.

[12] Miklós Ajtai, Vašek Chvátal, Monroe M. Newborn, and Endre Szemerédi. Crossing-free subgraphs. *Ann. Discrete Math.*, 12:9–12, 1982.

[13] Noga Alon and Shmuel Onn. Separable partitions. *Discrete Appl. Math.*, 91(1-3):39–51, 1999.

[14] Brian Alspach. Research problems. *Discrete Math.*, 36(3):333–334, 1981.

[15] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Appl. Math.*, 65(1-3):21–46, 1996.

[16] Imre Bárány and Zoltán Füredi. Empty simplices in Euclidean space. *Canad. Math. Bull.*, 30(4):436–445, 1987.

[17] Imre Bárány and Pavel Valtr. Planar point sets with a small number of empty convex polygons. *Studia Sci. Math. Hungar.*, 41(2):243–266, 2004.

[18] Harold William Becker. Planar rhyme schemes. *Bull. Amer. Math. Soc.*, 58:39, 1952.

[19] Nicla Bernasconi, Konstantinos Panagiotou, and Angelika Steger. On properties of random dissections and triangulations. In *SODA '08: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 132–141, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.

[20] Sergei Bespamyatnikh. An efficient algorithm for enumeration of triangulations. *Comput. Geom. Theory Appl.*, 23(3):271–279, 2002.

[21] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Comput. Geom. Theory Appl.*, 42(1):60–80, 2009.

[22] Kevin Buchin, Andreas Razen, Takeaki Uno, and Uli Wagner. Transforming spanning trees: A lower bound. *Comput. Geom. Theory Appl.*, 42(8):724–730, 2009. Special Issue on the 23rd European Workshop on Computational Geometry.

[23] Kevin Buchin and André Schulz. On the number of spanning trees a planar graph can have, unpublished manuscript.

[24] Sonja Čukić, Michael Hoffmann, Andreas Razen, and Tibor Szabó. Notes, 2006. 4th Gremo Workshop on Open Problems, Wislikofen.

[25] Karel Culik II and Derick Wood. A note on some tree similarity measures. *Inform. Process. Lett.*, 15(1):39–42, 1982.

[26] Markus Oswald Denny and Christian A. Sohler. Encoding a triangulation as a permutation of its point set. In *CCCG '97: Proceedings of the 9th Canadian Conference on Computational Geometry*, pages 39–43, 1997.

[27] Nachum Dershowitz and Shmuel Zaks. Enumerations of ordered trees. *Discrete Math.*, 31(1):9–28, 1980.

[28] Adrian Dumitrescu. Planar sets with few empty convex polygons. *Studia Sci. Math. Hungar.*, 36(1-2):93–109, 2000.

[29] Paul Henry Edelman. Chain enumeration and non-crossing partitions. *Discrete Math.*, 31(2):171–180, 1980.

[30] Paul Erdős, László Lovász, Gustavus J. Simmons, and Ernst Gabor Straus. Dissection graphs of planar point sets. In Jagdish Narain Srivastava et al., editor, *A Survey of Combinatorial Theory*, pages 139–149. North Holland, 1973.

[31] Philippe Flajolet and Marc Noy. Analytic combinatorics of non-crossing configurations. *Discrete Math.*, 204(1-3):203–229, 1999.

[32] Jerôme Galtier, Ferran Hurtado, Marc Noy, Stéphane Pérennes, and Jorge Urrutia. Simultaneous edge flipping in triangulations. *Internat. J. Comput. Geom. Appl.*, 13(2):113–133, 2003.

[33] Alfredo García, Marc Noy, and Javier Tejel. Lower bounds on the number of crossing-free subgraphs of K_N. *Comput. Geom. Theory Appl.*, 16(4):211–221, 2000.

[34] Sabine Hanke, Thomas Ottmann, and Sven Schuierer. The edge-flipping distance of triangulations. *J.UCS*, 2(8):570–579, 1996.

[35] Edward Frank Harding. The number of partitions of a set of N points in k dimensions induced by hyperplanes. *Proc. Edinburgh Math. Soc.*, 15(II):285–289, 1967.

[36] Carmen Hernando, Ferran Hurtado, and Marc Noy. Graphs of non-crossing perfect matchings. *Graphs Combin.*, 18(3):517–532, 2002.

[37] María Carmen Hernando, Ferran Hurtado, Alberto Márquez, Mercè Mora, and Marc Noy. Geometric tree graphs of points in convex position. *Discrete Appl. Math.*, 93(1):51–66, 1999.

[38] Michael Edward Houle, Ferran Hurtado, Marc Noy, and Eduardo Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs Combin.*, 21(3):325–331, 2005.

[39] Ferran Hurtado and Marc Noy. Counting triangulations of almost-convex polygons. *Ars Combin.*, 45:169–179, 1997.

[40] Ferran Hurtado, Marc Noy, and Jorge Urrutia. Flipping edges in triangulations. In *SCG '96: Proceedings of the Twelfth Annual Symposium on Computational Geometry*, pages 214–223, New York, NY, USA, 1996. ACM.

[41] Mark Jerrum. *Counting, Sampling and Integrating: Algorithms and Complexity.* Lectures in Mathematics. Birkhäuser, 2003.

[42] Stasys Jukna. *Extremal combinatorics. With applications in computer science.* Texts Theoret. Comput. Sci. EATCS Ser. Springer-Verlag, Berlin, 2001.

[43] Pieter Willem Kasteleyn. Graph theory and crystal physics. In Frank Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967.

[44] Meir Katchalski and Amram Meir. On empty triangles determined by points in the plane. *Acta Math. Hungar.*, 51(3-4):323–328, 1988.

[45] Naoki Katoh and Shin-ichi Tanigawa. Fast enumeration algorithms for non-crossing geometric graphs. In *SCG '08: Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, pages 328–337, New York, NY, USA, 2008. ACM.

[46] Gustav Robert Kirchhoff. Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Verteilung galvanischer Ströme geführt wird. *Ann. Phys. Chem.*, 72:497–508, 1847.

[47] Martin Klazar. On trees and noncrossing partitions. *Discrete Appl. Math.*, 82(1-3):263–269, 1998.

[48] Germain Kreweras. Une famille d'identités mettant en jeu toutes les partitions d'un ensemble fini de variables en un nombre donné de classes. *C. R. Acad. Sci. Paris Sér. A-B*, 270:1140–1143, 1970.

[49] Germain Kreweras. Sur les partitions non croisées d'un cycle. *Discrete Math.*, 1(4):333–350, 1972.

[50] Charles Lawrence Lawson. Transforming triangulations. *Discrete Math.*, 3:365–372, 1972.

[51] Charles Lawrence Lawson. Software for C1 surface interpolation. In John Richard Rice, editor, *Mathematical Software III*, pages 161–194, New York, 1977. Academic Press.

[52] Der-Tsai Lee and Arthur Kuo-In Lin. Generalized Delaunay triangulation for planar graphs. *Discrete Comput. Geom.*, 1(1):201–217, 1986.

[53] Sheng-Chyang Liaw, Hong-Gwa Yeh, Frank Kwang-Ming Hwang, and Gerard Jennhwa Chang. A simple and direct derivation for the number of noncrossing partitions. *Proc. Amer. Math. Soc.*, 126(6):1579–1581, 1998.

[54] László Lovász. On covering of graphs. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 231–236. Academic Press, New York, 1968.

[55] László Lovász. On the number of halving lines. *Ann. Univ. Sci. Budapest, Eötvös, Sect. Math.*, 14:107–108, 1971.

[56] Jiří Matoušek. *Lectures on discrete geometry*, volume 212 of *Grad. Texts in Math.* Springer-Verlag, New York, 2002.

[57] Paul McCabe. Lower bounding the number of straight-edge triangulations of planar point sets. Master's thesis, Universität des Saarlandes, 2003.

[58] Paul McCabe and Raimund Seidel. New lower bounds for the number of straight-edge triangulations of a planar point set. In *Proceedings of the 20th European Workshop on Computational Geometry*, pages 175–176, 2004.

[59] John W. Moon. *Counting labelled trees*, volume 1 of *Canadian Mathematical Monographs*. Canadian Mathematical Congress, Montreal, 1970.

[60] Theodore Samuel Motzkin. Relations between hypersurface cross ratios, and a combinatorial formula for partitions of a polygon, for permanent preponderance, and for non-associative products. *Bull. Amer. Math. Soc.*, 54(4):352–360, 1948.

[61] Monroe Newborn and William Oscar Jules Moser. Optimal crossing-free Hamiltonian circuit drawings of K_n. *J. Combin. Theory Ser. B*, 29(1):13–26, 1980.

[62] Helmut Prodinger. A correspondence between ordered trees and non-crossing partitions. *Discrete Math.*, 46(2):205–206, 1983.

[63] Andreas Razen. A lower bound for the transformation of compatible perfect matchings. In *Proceedings of the 24th European Workshop on Computational Geometry*, pages 115–118, 2008.

[64] Andreas Razen, Jack Snoeyink, and Emo Welzl. Number of crossing-free geometric graphs vs. triangulations. *Electronic Notes in Discrete Mathematics*, 31:195–200, 2008.

[65] Andreas Razen and Emo Welzl. On the number of crossing-free partitions. In *Proceedings of the 25th European Workshop on Computational Geometry*, pages 147–150, 2009.

[66] Andreas Razen and Emo Welzl. Counting crossing-free geometric graphs with exponential speed-up, submitted.

[67] Gerhard Ringel. Extremal problems in the theory of graphs. In *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*, pages 85–90. Publ. House Czechoslovak Acad. Sci., Prague, 1964.

[68] Günter Rote, Ares Ribó, and Xuerong Yong. The number of spanning trees in a planar graph. In *Oberwolfach Reports*, volume 2, pages 969–973. European Mathematical Society - Publishing House, 2005.

[69] Francisco Santos and Raimund Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Comb. Theory Ser. A*, 102(1):186–193, 2003.

[70] Raimund Seidel. On the number of triangulations of planar point sets. *Combinatorica*, 18(2):297–299, 1998.

[71] Micha Sharir and Emo Welzl. On the number of crossing-free matchings, cycles, and partitions. *SIAM J. Comput.*, 36(3):695–720, 2006.

[72] Micha Sharir and Emo Welzl. Random triangulations of planar point sets. In *SCG '06: Proceedings of the Twenty-second Annual Symposium on Computational Geometry*, pages 273–281, New York, NY, USA, 2006. ACM.

[73] Rodica Simion. Noncrossing partitions. *Discrete Math.*, 217(1-3):367–409, 2000. Formal power series and algebraic combinatorics (Vienna, 1997).

[74] Daniel Dominic Sleator, Robert Endre Tarjan, and William Paul Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.

[75] Warren Douglas Smith. *Studies in computational geometry motivated by mesh generation.* PhD thesis, Princeton University, 1989.

[76] Richard Peter Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Stud. Adv. Math.* Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

[77] Richard Peter Stanley. Catalan addendum to [76], August 2009. `http://www-math.mit.edu/~rstan/ec/`.

[78] Paul Turán. Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[79] William Thomas Tutte. *Graph Theory*, volume 21 of *Encyclopedia Math. Appl.* Addison-Wesley, 1984.

[80] Pavel Valtr. On the minimum number of empty polygons in planar point sets. *Studia Sci. Math. Hungar.*, 30(1-2):155–163, 1995.

[81] Uli Wagner. k-sets and k-facets. In *Surveys on discrete and computational geometry*, volume 453 of *Contemp. Math.*, pages 443–513. Amer. Math. Soc., Providence, RI, 2008.

[82] Douglas Brent West. *Introduction to graph theory.* Prentice Hall Inc., Upper Saddle River, NJ, 1996.

Glossary

$\mathbb{1}_{[A]}$	indicator function with predicate A. 21
$[\alpha, \beta] \times \mathbb{R}$	closed vertical strip in $\mathbb{R}^2$ between x-coordinates α and β. 63
$\mathrm{cfp}(P)$	number of crossing-free partitions of a point set P. 18
$\mathrm{cfp}_k(P)$	number of crossing-free partitions of a point set P into k classes. 18
$e_l(Q)$	number of long edges in a crossing-free partition Q. 51
$e_s(Q)$	number of short edges in a crossing-free partition Q. 51
$\mathbb{E}[X]$	expectation of a random variable X. 13
$\mathit{fl}(T)$	number of flippable edges in a triangulation T. 92
Γ_n	set of n points in convex position. 18
$I^P(Q)$	subset of points in P that are contained inside the convex hull of Q. 21
$\ell(T)$	number of Lawson edges in a triangulation T. 94
$L(T)$	set of Lawson edges in a triangulation T. 94
$M = M(P)$	number of edges in any triangulation on P. 92, 118
$(n_1, n_2, \ldots, n_k)$-partition	partition of a set into k classes of size n_1, n_2,. . .,n_k. 21
$\mathit{nfl}(T)$	number of non-flippable edges in a triangulation T. 92
$O(f(n))$, $\Omega(f(n))$, $\Theta(f(n))$	Landau notation. 12
∂P	boundary of a point set P. 12
$\overline{P}$	closure of a point set P. 12
P°	interior of a point set P. 12

$\mathbb{P}[A]$	probability of an event A. 12
$\binom{P}{k}$	set of k-element subsets of P. 21
$\mathsf{pg}(P)$	number of crossing-free geometric graphs on P. 92, 118
$\mathcal{P}m(P)$	set of crossing-free perfect matchings of P. 80
S_n	single-chain on n points. 19
$\mathcal{S}t(P)$	set of crossing-free spanning trees of P. 60
$T^*(G)$	constrained Delaunay triangulation of a crossing-free geometric graph G. 94
$\mathcal{T}_{\mathrm{pm}}(P)$	transformation graph of compatible perfect matchings of P. 80
$\mathsf{tr}(P)$	number of triangulations on P. 93, 118
$\mathcal{T}r(P)$	set of triangulations on P. 93
$\mathcal{T}_{\mathrm{st}}(P)$	transformation graph of compatible spanning trees of P. 61
$u \rightarrow v$	vint u may be flipped down to vint v. 105
vint	*v*ertex-*in*-*t*riangulation. 104
$X(P)$	set of extreme points of P. 21

Curriculum Vitae

Andreas Razen

born May 7, 1982 in Feldkirch, Austria

1992 - 2000	Bundesgymnasium Feldkirch, Austria Degree: Matura
2000 - 2005	Studies of Mathematics at ETH Zurich, Switzerland Degree: Dipl. Math. ETH
since April 2005	Ph.D. student at ETH Zurich, Switzerland Field: Theoretical Computer Science