

POINTS, LINES, AND CIRCLES

SOME CONTRIBUTIONS TO COMBINATORIAL GEOMETRY

vorgelegt von
Dipl.-Ing.
Manfred Scheucher

Von der Fakultät II – Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften
– Dr. rer. nat. –

genehmigte Dissertation

Promotionsausschuss

Vorsitzender: Prof. Dr. Boris Springborn

Gutachter: Prof. Dr. Stefan Felsner


Prof. Dr. Günter Rote

Prof. Dr. Torsten Mütze

Prof. Dr. Pavel Valtr

Tag der wissenschaftlichen Aussprache: 15. August 2019

Berlin 2020



Zusammenfassung

In dieser Dissertation untersuchen wir Probleme aus den Gebieten Kombinatorik und Algorithmische Geometrie auf elementaren geometrischen Entitäten (Punkten, Linien und Kreisen).

Im ersten Teil betrachten wir sogenannte „Erdős–Szekeres type problems“: Ein klassisches Resultat von Erdős und Szekeres aus dem Jahr 1935 besagt, dass für jede natürliche Zahl k jede hinreichend große Punktmenge in allgemeiner Lage eine Teilmenge von k Punkten in konvexer Lage enthält – sogenannte „ k -Gons“. Wir untersuchen eine berühmte Variation: Ein „ k -Loch“ ist ein k -Gon mit der zusätzlichen Eigenschaft, dass kein weiterer Punkt in der konvexen Hülle des k -Lochs liegt. Diese Variante unterscheidet sich vom ursprünglichen Problem da beliebig große Punktmengen ohne 7-Löcher existieren. Neben der Existenz von k -Löchern wurde auch die Anzahl an k -Löchern in n -Punktmengen intensiv untersucht. Jede Punktmenge hat mindestens quadratisch viele 3-Löcher und 4-Löcher, und Punktmengen mit nur quadratisch vielen Löchern existieren. Zu Beginn meines Studiums waren für 5-Löcher und 6-Löcher nur lineare untere Schranken bekannt.

In dieser Dissertation präsentieren wir die erste superlineare untere Schranke für die Anzahl an 5-Löchern. Dies ist die erste asymptotische Verbesserung seit Harborths linearer Schranke aus dem Jahr 1978. Für unseren Beweis kombinieren wir traditionelle Beweismethoden mit Lemmas, welche mithilfe von Computerunterstützung verifiziert wurden. Des Weiteren entwickeln wir ein auf Boolescher Aussagenlogik basierendes Framework, welches die Untersuchung diverser kombinatorischer Eigenschaften von Punktmengen mittels SAT-Solver erlaubt.

Im zweiten Teil untersuchen wir Arrangements von Kreisen. Um ein besseres Verständnis derer Strukturen zu erhalten und um geometrischen Schwierigkeiten auszuweichen, betrachten wir „Arrangements von Pseudokreisen“ – eine Verallgemeinerung von Kreisarrangements, die von Grünbaum in den 1970er Jahren eingeführt wurde. Ein Pseudokreisarrangement ist eine Menge von einfachen geschlossenen Kurven auf der Sphäre oder in der Ebene mit der Eigenschaft, dass je zwei dieser Kurven disjunkt sind oder sich in genau zwei Punkten schneiden. In seinem Buch präsentierte Grünbaum die Vermutung, dass jedes Digon-freie Arrangement von n sich paarweise schneidenden Pseudokreisen mindestens $2n - 4$ Dreieckszellen enthält. Wir präsentieren Arrangements, die diese Vermutung widerlegen und geben neue Schranken für die Anzahl an Dreieckszellen für unterschiedliche Klassen von Arrangements an.

Des Weiteren untersuchen wir die „Kreisbarkeit“ von Arrangements: Offensichtlich ist jedes Kreisarrangement auch ein Pseudokreisarrangement – es ist jedoch schwer zu entscheiden, ob ein gegebenes Pseudokreisarrangement durch Kreise repräsentiert werden kann. Mit Hilfe eines Computerprogrammes konnten wir alle kombinatorisch unterschiedlichen Arrangements mit bis zu 7 Pseudokreisen enumerieren. Für Arrangements mit 5 Pseudokreisen und für Digon-freie Arrangements von 6 sich paarweise schneidenden Pseudokreisen geben wir eine vollständige Klassifizierung an: Entweder geben wir eine Kreisrepräsentation oder einen Nichtkreisbarkeitsbeweis an. Diese Beweise basieren auf Inzidenztheoremen (z.B. dem Satz von Miquel) und Deformationsargumenten, in welchen wir Kreise einer angenommenen Kreisrepräsentation in einer kontrollierten Weise ausweiten oder schrumpfen.

Im dritten und letzten Teil fassen wir die Ergebnisse aus Teil I und Teil II zusammen, und diskutieren weiterführende Fragen.

Abstract

In this dissertation we investigate some problems from the field of combinatorics and computational geometry which involve basic geometric entities (points, lines, and circles).

In the first part we look at Erdős–Szekeres type problems: The classical theorem by Erdős and Szekeres from 1935 asserts that, for every natural number k , every sufficiently large point set in general position contains a subset of k points in convex position – a so called “ k -gon”. We will investigate the famous variant of “ k -holes”, which are k -gons with the additional property that no other point lies in the convex hull of the k -hole. This variant differs from the original setting as there exist arbitrarily large point sets which do not contain 7-holes. Besides the existence of holes, also the number of k -holes in sets of n points have been studied intensively. It is well-known that point configurations contain at least quadratically many 3- and 4-holes, respectively, while point configurations with only quadratically many holes exist. Concerning 5- and 6-holes, the best published lower bounds were only linear at the time when I started my studies.

In this thesis we present the first superlinear lower bound on the number of 5-holes, which is the first asymptotic improvement since Harborth’s linear lower bound from 1978. For our proof we combine classical paper-and-pen proofs with lemmas that were proven using heavy computer assistance. We also develop a framework based on Boolean logic to investigate various combinatorial properties of point sets with the aid of SAT solvers.

In the second part we investigate arrangements of circles. Towards a better understanding of their structure and also to get rid of geometric difficulties, we look at the more general setting of “arrangements of pseudocircles” which was first introduced by Grünbaum in the 1970’s. An arrangement of pseudocircles is a collection of simple closed curves on the sphere or in the plane such that any two of the curves are either disjoint or intersect in exactly two points, where the two curves cross. In his book, Grünbaum conjectured that every digon-free arrangement of n pairwise intersecting pseudocircles contains at least $2n - 4$ triangular cells. We present arrangements to disprove this conjecture and give new bounds on the number of triangular cells for various classes of arrangements.

Furthermore, we study the “circularizability” of arrangements: it is clear that every arrangement of circles is an arrangement of pseudocircles, however, deciding whether an arrangement of pseudocircles is isomorphic to an arrangement of circles is computationally hard. Using a computer program, we have enumerated all combinatorially different arrangements of up to 7 pseudocircles. For the class of arrangements of 5 pseudocircles and for the class of digon-free intersecting arrangements of 6 pseudocircles, we give a complete classification: we either provide a circle representation or a non-circularizability proof. For these proofs we use incidence theorems like Miquel’s and arguments based on continuous deformation, where circles of an assumed circle representation grow or shrink in a controlled way.

In the third and last part we summarize the results from Part I and Part II, and discuss further questions.

Credits

Content of this dissertation appeared in the following publications:

- A Superlinear Lower Bound on the Number of 5-Holes (with Oswin Aichholzer, Martin Balko, Thomas Hackl, Jan Kynčl, Irene Parada, Pavel Valtr, and Birgit Vogtenhuber). An extended abstract appeared in the proceedings of the European Workshop on Computational Geometry (EuroCG2017) [ABH⁺17a]. The short version appeared in the proceedings of the International Symposium on Computational Geometry (SoCG 2017) [ABH⁺17b]. The full version will appear in the Journal of Combinatorial Theory, Series A, and is available at [ABH⁺19a].
The content of this publication appears in Chapter 3, and in the Introduction and Preliminaries of Part I (Chapters 1 and 2).
- On Disjoint Holes in Point Sets. An extended abstract appeared in the proceedings of the European Workshop on Computational Geometry (EuroCG2019) [Sch19a] and in the proceedings of the European Conference on Combinatorics, Graph Theory and Applications (EUROCOMB'19) [Sch19b]. The full version is available at [Sch18].
The content of this publication appears in Chapter 4, and in the Introduction and Preliminaries of Part I (Chapters 1 and 2).
- A Note On Universal Point Sets for Planar Graphs (with Hendrik Schrezenmaier and Raphael Steiner). An extended abstract appeared in the proceedings of the European Workshop on Computational Geometry (EuroCG2019) [SSS19a]. The short version is to appear in the proceedings of the Graph Drawing conference (GD2019). The full version is available at [SSS19b].
From this publication, only the parts about of the enumeration of order types appear in the Preliminaries of Part I (Chapter 2).
- Arrangements of Pseudocircles: On Circularizability (with Stefan Felsner). An extended abstract appeared in the proceedings of the European Workshop on Computational Geometry (EuroCG2018) [FS18a]. The short version appeared in the proceedings of the Graph Drawing conference (GD2018) [FS18b]. The full version will appear in the Ricky Pollack Memorial Issue of the journal Discrete & Computational Geometry [FS19a] and is available at [FS17a].
The content of this publication appears in Chapter 7, and in the Introduction and Preliminaries of Part II (Chapters 5 and 6).
- Arrangements of Pseudocircles: Triangles and Drawings (with Stefan Felsner). An extended abstract appeared in the proceedings of the European Workshop on Computational Geometry (EuroCG2017) [FS17b]. The short version appeared in the proceedings of the Graph Drawing conference (GD2017) [FS18c]. The full version will appear in the journal Discrete & Computational Geometry and is available at [FS19b].
The content of this publication appears in Chapter 8, and in the Introduction and Preliminaries of Part II (Chapters 5 and 6).

Acknowledgements

During my doctoral studies I spent some time in Graz and Budapest before I finally arrived in Berlin. Gratefully I acknowledge support from:

- ESF EUROCORES programme EuroGIGA – CRP ComPoSe, Austrian Science Fund (FWF): I648-N18, and the Austrian Science Fund (FWF): P23629-N18 – thanks to Thomas Hackl, Oswin Aichholzer, and Birgit Vogtenhuber for their support;
- ERC Advanced Research Grant no. 267165 (DISCONV) – thanks to Imre Bárány for inviting me to Alfréd Rényi Institute of Mathematics for 5 months and also to Pavel Valtr, for making Budapest a possibility and also for his support during the last years;
- DFG Grants FE 340/11-1 and FE 340/12-1 – my deepest gratitude to Stefan for his great support during the last years and, in particular, for making this position in Berlin possible.

My sincerest thanks to Torsten Mütze, Günter Rote, Pavel Valtr, and Boris Springborn for very valuable comments, which allowed me to further improve the quality of this thesis. I also thank the anonymous reviewers for carefully going through the relevant papers and for their valuable comments that helped to improve the quality of this thesis and the overall presentation.

I also want to thank my friends, colleagues, and coauthors – and last but not least – my parents, Wilfried and Ernestine, my brother Markus, and Helena for their great support.

Contents

Zusammenfassung	5
Abstract	7
Credits	9
Acknowledgements	11
Contents	13
I Erdős–Szekeres-Type Problems	17
1 Erdős–Szekeres-Type Problems on Point Sets	19
1.1 k -Holes	20
1.2 Number of Holes	21
1.3 Disjoint Holes	22
2 Preliminaries	25
2.1 Order Types of Point Sets	25
2.2 Abstract Order Types	26
2.3 Enumeration of Order Types	27
2.4 Point-Line Duality	28
3 A Superlinear Lower Bound on the Number of 5-Holes	31
3.1 Outline	31
3.2 Proof of Theorem 3.1	32
3.3 Proof of Theorem 3.2	33
3.4 Preliminaries for the proof of Theorem 3.3	34
3.5 Proof of Theorem 3.3	36
3.5.1 Sequences of a^* -wedges with at most two points of B	38
3.5.2 Computer-assisted results	45
3.5.3 Applications of the computer-assisted results	46
3.5.4 Extremal points of ℓ -critical sets	47
3.5.5 Finalizing the proof of Theorem 3.3	51
3.6 Necessity of Assumptions	51
3.6.1 Necessity of the assumptions in Theorem 3.3	52
3.6.2 Necessity of the assumptions in Lemmas 3.14 to 3.17	54
4 Finding Holes using SAT Solvers	55
4.1 Three Disjoint Holes	56
4.2 Many Disjoint Holes	57

4.3	Encoding with Triple Orientations	59
4.3.1	Signotope Axioms	59
4.3.2	Increasing Coordinates and Cyclic Order	59
4.4	SAT Model	60
4.4.1	A Detailed Description	60
4.4.2	Unsatisfiability and Verification	62
4.5	Further Applications of the SAT Model	63
II	Arrangements of Pseudocircles	65
5	Arrangements of Pseudocircles	67
5.1	Circularizability	69
5.2	Triangular Cells	71
6	Preliminaries	73
6.1	Arrangements of Great-Pseudocircles	73
6.2	Incidence Theorems	75
6.3	Flips and Deformations	76
7	On Circularizability	79
7.1	The Great-Circle Theorem and its Applications	81
7.2	Arrangements of 5 Pseudocircles	83
7.2.1	Non-circularizability of \mathcal{N}_5^1	83
7.2.2	Non-circularizability of the connected arrangements \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4	85
7.3	Intersecting Digon-free Arrangements of 6 Pseudocircles	86
7.3.1	Non-circularizability of \mathcal{N}_6^Δ	86
7.3.2	Non-circularizability of \mathcal{N}_6^2	89
7.3.3	Non-circularizability of \mathcal{N}_6^3	91
7.4	Additional Arrangements with $n = 6$	92
7.4.1	Non-circularizability of 3 intersecting arrangements with $n = 6$	92
7.4.2	Non-circularizability of 5 connected digon-free arrangements with $n = 6$	95
7.5	Enumeration and Asymptotics	99
7.5.1	Enumeration of Arrangements	100
7.5.2	Generating Circle Representations	100
7.5.3	The Circularizability Problem as System of Polynomial Inequalities	101
7.5.4	Counting Arrangements	102
7.5.5	Connectivity of the Flip-Graph	103
8	Triangles in Arrangements of Pseudocircles	105
8.1	Intersecting Arrangements with few Triangles	105
8.1.1	Intersecting Arrangements with Digons	109
8.2	Maximum Number of Triangles	111
8.2.1	Constructions using Arrangements of Pseudolines	115
III	Discussion	117
9	Discussion	119
9.1	Holes in Point Sets	119
9.2	Disjoint Holes in Point Sets	120
9.3	Erdős-Szekeres type Questions on Colored Point Sets	120
9.4	Triangles in Arrangements of Pseudocircles	121

9.5	Digons in Arrangements of Pseudocircles	122
9.6	Larger Cells in Arrangements	123
9.7	Circularizability	123
9.8	Number of Arrangements	124
9.9	Convexibility	124
9.10	Extending Arrangements	125
9.11	Cylindrical Arrangements	126
9.12	Flipgraph of Arrangements	126
9.13	Colorability of Arrangements	127
9.14	A Generalization of the Erdős–Szekeres Theorem	127
9.15	Encoding of Arrangements	127
9.16	Testing Isomorphism of Arrangements	128
9.17	SAT-Encoding of Arrangements	128
9.18	A Generalization of the Zone Theorem	128
9.19	Visualization of Arrangements	129
Bibliography		131

Part I

Erdős–Szekeres-Type Problems

Chapter 1

Erdős–Szekeres-Type Problems on Point Sets

A set of points in the Euclidean plane $S \subseteq \mathbb{R}^2$ is *in general position* if no three points lie on a common line. Throughout this thesis all point sets are considered to be finite and in general position unless mentioned otherwise. A subset $X \subseteq S$ of size $|X| = k$ is a *k-gon* if all points of X lie on the boundary of the convex hull of X . Figure 1.1 gives an illustration.



Figure 1.1: An illustration of a 5-gon (left) and a 6-gon (right).

A classical result in combinatorial geometry and Ramsey theory from the 1930s by Erdős and Szekeres asserts that, for fixed $k \in \mathbb{N}$, every sufficiently large point set contains a k -gon [ES35].

Theorem ([ES35], The Erdős–Szekeres Theorem). *For every integer $k \geq 3$, there is a smallest integer $g(k)$ such that every set of at least $g(k)$ points in general position in the plane contains k points in convex position.*

Paul Erdős named the problem of determining the value $g(k)$ the “Happy ending problem” because George Szekeres and Esther Klein became engaged while working on the problem and got married subsequently (cf. [Hof98]).

In their paper [ES35], Erdős and Szekeres provided two proofs of the theorem (see also Chapter 3.1 of Matoušek’s book [Mat02]). In the first proof, they color each 4-tuple of points red if it is in convex position, and blue otherwise. It then follows from Ramsey–Theory¹ that there is a large monochromatic clique, and since blue 5-cliques cannot exist, there is a large red clique. Carathéodory’s theorem then asserts that all points from this clique are in convex position. In their second proof, they made more elaborate use of the relative positions of points to show the upper bound $g(k) \leq \binom{2k-4}{k-2} + 1$. Using Stirling’s formula, it is easy to see that this bound is of order $4^{k-o(k)}$. There were several small improvements on the upper bound by various researchers in the last decades, each of order $4^{k-o(k)}$ [CG98, KP98, TV98, TV05, MV16, NY16], until Suk [Suk17] significantly improved the upper bound to $2^{k+O(k^{2/3} \log k)}$. This bound was further improved to $2^{k+O(\sqrt{k} \log k)}$ by Holmsen and others [HMPT17].

¹We refer the reader who is not yet familiar with Ramsey–Theory for hypergraphs to a standard textbook such as “Ramsey Theory” by Graham, Rothschild, and Spencer [GRS90].

Concerning the lower bound on $g(k)$, Erdős and Szekeres constructed point sets of size 2^{k-2} with no k -gons, which they conjectured to be the largest possible sets with no k -gons. In fact, Erdős offered \$500 for a proof of the conjecture in the 1990's shortly before his death. Today $g(k) = 2^{k-2} + 1$ is confirmed only for up to $k = 6$, the values of $g(k)$ for $k \geq 7$, however, remain unknown. To be more precise, the value $g(4) = 5$ was first observed by Klein, $g(5) = 9$ by Kalbfleisch, Kalbfleisch, and Stanton [KKS70], and the value $g(6) = 17$ was determined by Szekeres and Peters using heavy computer assistance [SP06]. In the course of Chapter 4 we will come across a SAT model which allows us to verify $g(6) = 17$ with significantly smaller computation time than the original program by Szekeres and Peters [SP06].

1.1 k -Holes

The Erdős-Szekeres Theorem motivated a lot of further research, including numerous modifications and extensions of the theorem². In the 1970s, Erdős [Erd78] asked whether every sufficiently large point set contains a k -hole, that is, a k -gon with no other points of S lying inside its convex hull. Figure 1.2 gives an illustration. Analogous to $g(k)$, we define $h(k)$ as the smallest number (if such a number exists) such every set of at least $h(k)$ points in the plane contains a k -hole. It is worth mentioning that also “non-convex holes” have been investigated in literature (cf. Section 9.3 in the Discussion), however, throughout this dissertation we consider all holes to be convex.



Figure 1.2: An illustration of a 4-gon which is not a 4-hole (left) and a 4-hole (right).

It is not hard to see that $h(3) = 3$ and $h(4) = 5$. Harborth [Har78] proved by an elaborate case distinction that there is a 5-hole in every set of 10 points in general position in the plane and gave a construction of 9 points in general position with no 5-hole; hence $h(5) = 10$. After unsuccessful attempts of researchers to answer Erdős' question affirmatively for any fixed integer $k \geq 6$, Horton [Hor83] constructed, for every positive integer n , a set of n points in general position in the plane with no 7-hole (see also Chapter 3.2 of Matoušek's book [Mat02]). Horton's construction – which is today known as the *perfect Horton set* – was later generalized to so-called *Horton sets* and *squared Horton sets* [Val92a], and also to higher dimensions [Val92b].

The question whether there is a 6-hole in every sufficiently large finite planar point set remained open until 2007, when Gerken [Ger08] and Nicolás [Nic07] independently gave an affirmative answer. In both proofs the key idea is that sets with sufficiently large gons contain 6-holes. A reasonably short proof for the existence of 6-holes was later given by Valtr [Val08], who showed $h(6) \leq g(216)$. The currently best bound is by Koshelev [Kos09b]³, who showed that every set of 463 points contains a 6-hole. However, the largest set without 6-holes currently known has 29 points and was found using computer-assistance by Overmars [Ove02]. The gap between the upper and the lower bound on $h(6)$ remains huge.

²We refer the interested reader to the survey of Morris and Soltan [MS00]

³Koshelev's publication covers more than 50 pages (written in Russian)

1.2 Number of Holes

For positive integers n and k , let $h_k(n)$ be the minimum number of k -holes in a set of n points in the plane. From perfect Horton sets it clearly follows that $h_k(n) = 0$ holds for every n and every $k \geq 7$.

3-Holes and 4-Holes The functions $h_3(n)$ and $h_4(n)$ are both known to be asymptotically quadratic. A quadratic lower bound on $h_3(n)$ can be obtained easily from the fact that every pair of points a, b , together with the point c closest to the line \overline{ab} , spans a 3-hole. A similar idea was applied to 4-holes by Bárány and Füredi [BF87]: For every crossed line-segment ab one can find a 4-hole which has this segment as a diagonal. Since the non-crossed line-segments of a point set form a planar graph – which only has linearly many edges – this gives the quadratic lower bound $h_4(n) \geq \frac{n^2}{4} - O(n)$. Bárány and Füredi further showed $h_3(n) \geq n^2 - O(n \log n)$, that randomly chosen point sets contain only quadratically many 3-holes (see also [Val95]), and that perfect Horton sets have only quadratically many 4-holes, 5-holes, and 6-holes.

After a couple of subsequent improvements [Deh87, KM88, Val95, Dum00, Gar12], the currently best known lower bounds today are due to Aichholzer et al. [AFMH⁺14] and the best known upper bounds are due to Bárány and Valtr [BV04]:

$$\begin{aligned} n^2 - \frac{32n}{7} + \frac{22}{7} &\leq h_3(n) \leq 1.6196n^2 + o(n^2) \\ \frac{n^2}{2} - \frac{9n}{4} - o(n) &\leq h_4(n) \leq 1.9397n^2 + o(n^2). \end{aligned}$$

It is worth mentioning that the upper bounds in [BV04] were obtained for the above mentioned squared Horton sets [Val92a], which are basically $\sqrt{n} \times \sqrt{n}$ grids perturbed in a way such that all collinear points become Horton sets.

5-Holes and 6-Holes For $h_5(n)$ and $h_6(n)$, no matching bounds are known. Both functions, however, are conjectured to be asymptotically quadratic. So far, the best known upper bounds $h_5(n) \leq 1.0207n^2 + o(n^2)$ and $h_6(n) \leq 0.2006n^2 + o(n^2)$ were obtained by Bárány and Valtr [BV04].

As noted by Bárány and Füredi [BF87], a linear lower bound of $\lfloor n/10 \rfloor$ follows directly from Harborth's result [Har78]. Bárány and Károlyi [BK01] improved this bound to $h_5(n) \geq n/6 - O(1)$. In 1987, Dehnhardt [Deh87] showed $h_5(11) = 2$ and $h_5(12) = 3$ in his PhD thesis⁴. His result, however, remained unknown to the scientific community for a long time, until García [Gar12] used the values obtained by Dehnhardt to prove the lower bound $h_5(n) \geq 3\lfloor \frac{n-4}{8} \rfloor$. Further improvement followed by Aichholzer, Hackl, and Vogtenhuber [AHV12] and by Valtr [Val12]. The best known bound published before my studies was by Aichholzer et al. [AFMH⁺14], who showed $h_5(n) \geq 3n/4 - o(n)$. All above mentioned improvements on the multiplicative constant were achieved by utilizing the values of $h_5(10)$, $h_5(11)$, and $h_5(12)$. The lower bound on $h_6(n)$ has a similar (but shorter) history. The best bound currently known is by Valtr [Val12], who showed $h_6(n) \geq n/229 - 4$.

In my bachelor's thesis [Sch13] the exact values $h_5(13) = 3$, $h_5(14) = 6$, and $h_5(15) = 9$ were determined and $h_5(16) \in \{10, 11\}$ was shown. During the preparation of this thesis, we further determined the value $h_5(16) = 11$ using a SAT model, which we will describe later in Chapter 4.5. Table 1.1 summarizes our knowledge on the values of $h_5(n)$ for $n \leq 20$ (cf. sequence A276096 on the OEIS [Slo]). The values $h_5(n)$ for $n \leq 16$ can be used to obtain further improvements on the multiplicative constant. By revising the proofs of [AFMH⁺14, Lemma 1] and [AFMH⁺14, Theorem 3], one can obtain $h_5(n) \geq n - 10$ and $h_5(n) \geq 3n/2 - o(n)$, respectively.

⁴Harborth was Dehnhardt's supervisor.

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

Table 1.1: The minimum number $h_5(n)$ of 5-holes determined by any set of $n \leq 20$ points.

Relations Edelman and Jamison [EJ85] and Ahrens, Gordon, and McMahon [AGM99] have shown that every set P of n points fulfills the two relations

$$\begin{aligned} h_1(P) - h_2(P) + h_3(P) - h_4(P) \pm \dots &= 1 \\ h_1(P) - 2h_2(P) + 3h_3(P) - 4h_4(P) \pm \dots &= i(P), \end{aligned}$$

where $h_1(P) = \binom{n}{1}$, $h_2(P) = \binom{n}{2}$, $h_k(P)$ for $k \geq 3$ denotes the number of k -holes in P , and $i(P)$ denotes the number of inner points of P . Later Edelman and Reiner [ER00] gave an independent proof for the two relations (cf. [PRS06] and [BV04]). Those relations turned out to be quite powerful when working with k -holes. For example, for any point set P without 7-holes only the four terms $h_k(P)$, $3 \leq k \leq 6$, depend on the structure of P , and hence, any three of those terms determine the fourth one. Two further relations between k -holes have been shown by Pinchasi, Radoičić, and Sharir [PRS06]:

$$h_4(P) \geq h_3(P) - \frac{1}{2}n^2 - O(n) \quad \text{and} \quad h_5(P) \geq h_3(P) - n^2 - O(n).$$

Consequently, $h_3(n) \geq (1 + \epsilon)n^2 - o(n^2)$ would imply a quadratic lower bound on $h_5(n)$, or conversely, a subquadratic upper bound on $h_5(n)$ would imply that $h_3(n) \leq n^2 + o(n^2)$. Here is also worth mentioning that two further relations were shown by García [Gar12]. However, since they require some additional definitions, those two relations are deferred to Chapter 3 (cf. Theorem 3.5).

In Chapter 3, we give the first superlinear lower bound on the number of 5-holes, which is the first asymptotic improvement since Harborth's linear lower bound from 1978. In particular, we show $h_5(n) \geq \Omega(n \log^{4/5} n)$. This is one of the main results of this thesis and solves an open problem, which was explicitly stated, for example, in a book by Brass, Moser, and Pach [BMP05, Chapter 8.4, Problem 5] and in the survey by Aichholzer [Aic09]. Moreover, we slightly adapt our proof and provide improved lower bounds on the minimum numbers of 3-holes and 4-holes: $h_3(n) \geq n^2 + \Omega(n \log^{2/3} n)$ and $h_4(n) \geq \frac{n^2}{2} + \Omega(n \log^{3/4} n)$.

1.3 Disjoint Holes

In 2001, Hosono and Urabe [HU01] started the investigation of disjoint holes, where two holes X_1, X_2 of a given point set S are said to be *disjoint* if their respective convex hulls are disjoint (that is, $\text{conv}(X_1) \cap \text{conv}(X_2) = \emptyset$); see Figure 1.3 for an illustration. This led to the following question: What is the smallest number $h(k_1, \dots, k_l)$ such that every set of $h(k_1, \dots, k_l)$ points determines a k_i -hole for every $i = 1, \dots, l$, such that the holes are pairwise disjoint [HU08]?




Figure 1.3: An illustration of (interior-)disjoint 3-holes.

As there are arbitrarily large point sets without 7-holes, only parameters $k_i < 7$ are of interest. Moreover, since the gap between the upper bound and the lower bound for $h(6)$ is still huge, mostly values with parameters $k_1, \dots, k_l \leq 5$ were investigated. Also note that, if all k_i are at most 3, then the value $h(k_1, \dots, k_l) = k_1 + \dots + k_l$ is straight-forward because every set of $k_1 + \dots + k_l$ points can be cut into blocks of k_1, \dots, k_l points (from left to right), which clearly determine the desired holes.

For two parameters, the value $h(k_1, k_2)$ has been determined for all $k_1, k_2 \leq 5$ except for $h(5, 5)$ [HU01, HU05, HU08, BD11]. Concerning the value $h(5, 5)$, the previous best bounds are $17 \leq h(5, 5) \leq 19$. In particular, the upper bound was shown by Bhattacharya and Das [BD13] by an elaborate case distinction.

In Chapter 4, we present a technique, based on a SAT model, which can be used to tackle various combinatorial problems on (relatively small) point sets. With this approach, we show that every set of 17 points in general position admits two disjoint 5-holes, that is, $h(5, 5) = 17$. This answers a question of Hosono and Urabe [HU01]. We also provide new bounds for three and more pairwise disjoint holes. Table 1.2 shows the best possible bounds for two disjoint holes.

	2	3	4	5
2	4	5	6	10
3		6	7	10
4			9	12
5				17*

Table 1.2: Values of $h(k_1, k_2)$. The entry marked with star (*) is new.

Interior-Disjoint Holes Besides disjoint holes, also the variant of interior-disjoint holes has been investigated intensively by various groups of researchers (see e.g. [DHKS03, SU07, CGH⁺15, BMS17, HU18]). Interior-disjoint holes also play an important role in the study of other geometric objects such as visibility graphs (see e.g. [DTP09]) or flip graphs of triangulations on point sets (see e.g. [Pil18]). Wagner and Welzl [WW19] quite recently developed a framework for triangulations on planar point sets which also allows the investigation of interior-disjoint holes. Using their tools and results it is, for example, quite easy to derive that 10 points always give a 4-hole and a 5-hole that are interior-disjoint. In fact, when comparing the number of researchers working on the respective problems, the interior-disjoint holes appear to be more of interest.

Two holes X_1, X_2 are called *interior-disjoint* if their respective convex hulls are interior-disjoint; see Figure 1.3. Interior-disjoint holes are also called *compatible holes* in literature. Note that a pair of interior-disjoint holes can share up to two vertices. In a recent article from 2018, Hosono and Urabe [HU18] summarized the current status and presented some new results.

Using our SAT model from Chapter 4, we show that every set of 15 points contains two interior-disjoint 5-holes – this strengthens the result from Hosono and Urabe [HU18]. Table 1.3 shows the best possible bounds for two interior-disjoint holes.

	3	4	5
3	4	5	10
4		7	10
5			15*

Table 1.3: Best possible bounds on the minimum number of points such that every set of that many points contains two interior-disjoint holes of sizes k_1 and k_2 . The entry marked with star (*) is new.

Chapter 2

Preliminaries

2.1 Order Types of Point Sets

Even though, for fixed $n \in \mathbb{N}$, there is an uncountable number of possibilities to place n points in the Euclidean plane, there are only finitely many equivalence classes of point sets when point sets inducing the same triple orientations are considered equal. Given a set of points $S = \{s_1, \dots, s_n\}$ with $s_i = (x_i, y_i)$, we say that the triple (a, b, c) is *positively (negatively) oriented* if

$$\chi_{abc} := \text{sgn} \det \begin{pmatrix} 1 & 1 & 1 \\ x_a & x_b & x_c \\ y_a & y_b & y_c \end{pmatrix} \in \{-1, 0, +1\}$$

is positive (negative)¹. As illustrated in Figure 2.1, the point c lies to the left of the directed line \overrightarrow{ab} if and only if the triple (a, b, c) is positively oriented. Also note that $\chi_{abc} = 0$ indicates collinear points, in particular, $\chi_{aaa} = \chi_{aab} = \chi_{aba} = \chi_{baa} = 0$. As introduced by Goodman and Pollack [GP83], these equivalence classes (sometimes also with unlabeled points) are called *order types*.




Figure 2.1: An illustration of orientation triples: (a, b, c) is positively and (a, b, d) is negatively oriented.

It is easy to see that many properties of point sets such as convexity are rather combinatorial than geometric properties as they can be described solely using the relative position of the points: If the points s_1, \dots, s_k are the vertices of a convex polygon (ordered along the boundary), then, for every $i = 1, \dots, k$, the cyclic order of the other points around s_i is $s_{i+1}, s_{i+2}, \dots, s_{i-1}$ (indices modulo k). Equivalently, for every $i = 1, \dots, k$, the points s_{i+2}, s_{i+3}, \dots lie on one side of the line $\overline{s_i s_{i+1}}$. Similarly, one can also describe k -holes only using relative positions: A point s_0 lies inside a triangle spanned by points s_a, s_b, s_c if

- (i) the line $\overline{s_0 s_a}$ separates s_b from s_c ,
- (ii) the line $\overline{s_0 s_b}$ separates s_a from s_c , and
- (iii) the line $\overline{s_0 s_c}$ separates s_a from s_b .

¹The letter χ is commonly used in literature to denote triple orientations as the word “chirality” is derived from the Greek word for “hand”.

Carathéodory's theorem now asserts that s_0 lies inside a convex polygon spanned by the points s_1, \dots, s_k if and only if there is a triangle s_a, s_b, s_c ($1 \leq a < b < c \leq k$) containing s_0 .

In the course of this thesis we will come across point sets, that are separated by a line. Also this property only depends on the order type: Suppose that a line ℓ separates point sets A and B . Then, for example by rotating ℓ , we can find another line ℓ' that contains a point $a \in A$ and a point $b \in B$ and separates $A \setminus \{a\}$ and $B \setminus \{b\}$. In particular, we have $\chi_{aba'} \leq 0$ for all $a' \in A$ and $\chi_{abb'} \geq 0$ for all $b' \in B$, or the other way round.

2.2 Abstract Order Types

In the course of this thesis, we will use computer assistance to prove a couple of statements for (relatively small) point sets. For such computer proofs it is common practice to consider a generalization of order types – the so called *abstract order types* – which are solely described by triple orientations. In the first glance, one might think that it is a good idea to use a computer to simply test all possible mappings $\chi: \{1, \dots, n\}^3 \rightarrow \{-1, 0, +1\}$. This approach, however, might not result in an accurate result as only very few of such mappings describe actual point sets². Before we can properly define abstract order types, we recall some basics from projective geometry and linear algebra.

A point (x, y) in the plane can be identified with the line $\{(\lambda, \lambda x, \lambda y) : \lambda \in \mathbb{R}\}$ in 3-space, which in particular contains the origin. Obviously, every point (λ, x, y) in 3-space with $\lambda \neq 0$ describes the unique point $(x/\lambda, y/\lambda)$.

The well-known Graßmann–Plücker relations (see e.g. [BLW⁺99, Chapter 3.5]) assert that any r -dimensional vectors $a_1, \dots, a_r, b_1, \dots, b_r$ fulfill³

$$\det(a_1, \dots, a_r) \cdot \det(b_1, \dots, b_r) = \sum_{i=1}^r \det(b_i, a_2, \dots, a_r) \cdot \det(b_1, \dots, b_{i-1}, a_1, b_{i+1}, \dots, b_r).$$

In particular, two sides of the equation clearly must have the same sign. This necessary condition directly translates to point sets and their induced triple orientations. Moreover, from the properties of the determinant we also derive

$$\det(a_{\sigma(1)}, \dots, a_{\sigma(r)}) = \text{sgn}(\sigma) \cdot \det(a_1, \dots, a_r)$$

for any permutation σ of the indices $\{1, \dots, r\}$. These two necessary conditions clearly motivate the following generalization of point sets.

Definition 2.1 (Chirotope). *A function $\chi: \{1, \dots, n\}^r \rightarrow \{-1, 0, +1\}$ is called chirotope of rank r if the following three properties are fulfilled:*

- (i) χ not identically zero;
- (ii) for every permutation σ and indices $a_1, \dots, a_r \in \{1, \dots, n\}$,

$$\chi(a_{\sigma(1)}, \dots, a_{\sigma(r)}) = \text{sgn}(\sigma) \cdot \chi(a_1, \dots, a_r);$$

²In fact, mappings $\chi: \binom{\{1, \dots, n\}}{3} \rightarrow \{-1, +1\}$, which can be considered as bicolored 3-uniform hypergraphs, also have been studied as a natural generalization of point sets in terms of the Erdős–Szekeres Theorem (see e.g. [SP06, FPSS12, BV17])

³The relations can be derived for example as outlined: Consider the vectors $a_1, \dots, a_r, b_1, \dots, b_r$ as an $r \times 2r$ matrix and apply row additions to obtain echelon form. If the first r columns form a singular matrix, then $\det(a_1, \dots, a_r) = 0$ and both sides of the equation vanish by a simple column multiplication argument. Otherwise, we can assume (also due to a column multiplication argument) that the first r columns form an identity matrix. Since the determinant is invariant to row additions, none of the terms in the Graßmann–Plücker relations is effected during the transformation, and the statement then follows from Laplace expansion.

(iii) for indices $a_1, \dots, a_r, b_1, \dots, b_r \in \{1, \dots, n\}$,

$$\text{if } \chi(b_i, a_2, \dots, a_r) \cdot \chi(b_1, \dots, b_{i-1}, a_1, b_{i+1}, \dots, b_r) \geq 0 \text{ holds for all } i = 1, \dots, r \\ \text{then we have } \chi(a_1, \dots, a_r) \cdot \chi(b_1, \dots, b_r) \geq 0.$$

Chirotopes of rank 3 were also investigated under several different names (also with slight variations): *abstract order types* [AAK02, Kra03], *CC systems* (see [Knu92]), *pseudoconfigurations of points*, *abstract oriented matroids of rank 3* (see e.g. [BLW⁺99]), or *pseudolinear drawings of the complete graph K_n* (see e.g. [AMRS18]). Since we are mainly interested in planar point sets in general position in this thesis, we assume that abstract order types are *in general position* unless mentioned otherwise, that is, $\chi(a, b, c) \neq 0$ holds for any three distinct indices a, b, c . Abstract order types which are not in general position are sometimes also called *degenerate*.

Since property (iii) from Definition 2.1 – commonly known as the *exchange property* – can be considered as a relaxation of the Graßmann–Plücker relations, it is not that surprising that there are in fact abstract order types which are **not** induced by any point set – but we will defer this to Part II of this thesis (cf. Chapters 6.2 and 7).

2.3 Enumeration of Order Types

Given an abstract order type one can iteratively remove points from the convex hull until all points are removed. (Similar as in point sets, a and b are consecutive on the boundary of the convex hull if $\chi(a, b, c) = +$ holds for all other indices c .) In reverse, one can recursively enumerate all abstract order types on n points by taking a chirotope on $n - 1$ points and extending it by an n -th point in all possible ways such that the chirotope axioms remain satisfied (cf. Definition 2.1).

Aichholzer, Aurenhammer, and Krasser [AAK02, AK06] (see also Krasser’s PhD thesis [Kra03]) provide the *order type database* which stores all order types encoded as point sets of size up to 11. In fact, they enumerated all abstract order types on $n \leq 11$ points and either provided a realizing point set or a non-realizability proof. The set of order types of sets of 10 points are available online [Aic] and the ones for $n = 11$ are available upon request from Aichholzer since the database needs about 96 GB of storage. Table 2.1 shows the number of abstract and realizable order types for $n \leq 11$ (cf. sequences A006247 and A063666 on the OEIS [Slo]).

number of points	3	4	5	6	7	8
abstract order types	1	2	3	16	135	3 315
realizable order types	1	2	3	16	135	3 315

number of points	9		10		11	
abstract order types	158 830		14 320 182		2 343 203 071	
realizable order types	158 817		14 309 547		2 334 512 907	

Table 2.1: The number of abstract and realizable order types on n points for $n \leq 11$. [AK06]

In the course of another recent paper with two of my colleagues Hendrik Schrezenmaier and Raphael Steiner [SSS19b], we wrote a short and simple C++ program that allows the recursive enumeration of all abstract order types on 11 points with only one day of computation time on a single CPU. In particular, with our independent program we could verify the results for abstract order types from Aichholzer, Aurenhammer, and Krasser. Since none of our computer proofs in this thesis relies on coordinates, it is sufficient for our proofs to test the database of abstract

order types. In fact, also some non-realizable abstract order types are enumerated, however, all statements actually hold in the more general setting of abstract order types.

2.4 Point-Line Duality

When considering a point set $S = \{s_1, \dots, s_n\}$ in the plane, we may further assume that the points s_1, \dots, s_n have increasing x -coordinates. Note that when considering point sets as representatives for order types this assumption is not a restriction. In fact, given a set of point where some points have the same x -coordinate, we can apply a slight rotation (which preserves the order type) to obtain distinct x -coordinates.

Using the *unit paraboloid duality transformation*, which maps a point $s = (a, b)$ to the line $s^* : y = 2ax - b$, we obtain the arrangement of dual lines $S^* = \{s_1^*, \dots, s_n^*\}$, where the dual lines s_1^*, \dots, s_n^* have increasing slopes. By the increasing x -coordinates and the properties of the unit paraboloid duality (cf. [O'R94, Chapter 6.5] or [Ede87, Chapter 1.4]), the following three statements are equivalent for any indices $1 \leq i < j < k \leq n$:

- (i) The points s_i, s_j, s_k are positively oriented.
- (ii) The point s_k lies above the line $\overline{s_i s_j}$.
- (iii) The intersection point of the two lines s_i^* and s_j^* lies above the line s_k^* .

As a consequence, we can read the triple-orientations of a point set S also from its dual line arrangement S^* ; see Figure 2.2. Pause to note that the choice of the duality is not that important: The above equivalences applies to any duality transformation $(a, b) \mapsto \{(x, y) : y = \alpha ax - \beta b\}$ with $\alpha \cdot \beta > 0$. For $\alpha \cdot \beta < 0$ the roles of “above” and “below” are simply exchanged in the dual.




Figure 2.2: Point-line duality illustration.

All combinatorial properties that only depend on the order type (not the explicit coordinates) can be read from the dual line arrangement. To give an example, observe that two points a and b of a set S are consecutive on the boundary of the convex hull of S if all other points of $S \setminus \{a, b\}$ lie on one side of the line \overline{ab} . This directly translates to the dual line arrangement: the intersection point of a^* and b^* either lies above all or below all lines from $S^* \setminus \{a^*, b^*\}$. Moreover, S is in convex position if and only if every line from S^* bounds the top or the bottom cell. Here the *top cell* and the *bottom cell* denote the two cells which contain the points $(0, +\infty)$ and $(0, -\infty)$, respectively.

Another way to define a point-line duality is as follows: Consider a collection L of lines in 3-space which pass through the origin. The intersection of these lines with a fixed plane Π yields a point set which lives in Π . Now consider the collection E of planes orthogonal to the lines from L and containing the origin. The intersection of the planes from E with another fixed plane Π' yields an arrangement of lines which lives in Π' . Figure 2.3⁴ gives an illustration.

⁴This 3D-figure was created using GeoGebra [H⁺19]




Figure 2.3: Point-line duality illustration in 3 dimensions. The black line L passes through the origin O and intersects Π in the red point P , and the black plane – which is orthogonal to L and contains the origin – intersects Π' in the blue line P^* .

Here it is also worth mentioning that by intersecting the planes from E with the unit sphere \mathbb{S}^2 (centered at the origin) one also gets a one-to-one correspondence to great-circles, however, we defer circle arrangements until Part II of this thesis (cf. Figure 6.1).

Chapter 3

A Superlinear Lower Bound on the Number of 5-Holes

In this chapter, we show $h_5(n) = \Omega(n \log^{4/5} n)$, obtaining the first superlinear lower bound on the minimum number of 5-holes among all sets of n points in the plane in general position.

Theorem 3.1. *There is an absolute constant $c > 0$ such that for every integer $n \geq 10$ we have $h_5(n) \geq cn \log^{4/5} n$.*

Moreover, using a result of García [Gar12], we adapt the proof of Theorem 3.1 to provide improved lower bounds on the minimum numbers of 3-holes and 4-holes.

Theorem 3.2. *The following two bounds are satisfied for every positive integer n :*

- (i) $h_3(n) \geq n^2 + \Omega(n \log^{2/3} n)$ and
- (ii) $h_4(n) \geq \frac{n^2}{2} + \Omega(n \log^{3/4} n)$.

3.1 Outline

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

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

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

The following structural result, which might be of independent interest, is a crucial step in the proof of Theorem 3.1.

Theorem 3.3. *Let $P = A \cup B$ be an ℓ -divided set with $|A|, |B| \geq 5$ and with neither A nor B in convex position. Then there is an ℓ -divided 5-hole in P .*

The proof of Theorem 3.3 is computer-assisted. We reduce the result to several statements about point sets of size at most 11 and then verify each of these statements by an exhaustive computer search. To verify the computer-aided proofs we have implemented two independent

programs, which, in addition, are based on different abstractions of point sets; see Chapter 3.5.2. Some of the tools that we use originate from my bachelor's theses [Sch13, Sch14].

Theorem 3.2 is proved in Chapter 3.3. Then, in Chapter 3.4, we give some preliminaries for the proof of Theorem 3.3, which is presented in Chapter 3.5. Finally, in Chapter 3.6, we give some final remarks. In particular, we show that the assumptions in Theorem 3.3 are necessary. To provide a better general view, we present a flow summary of the proof of Theorem 3.1 in Figure 3.1.

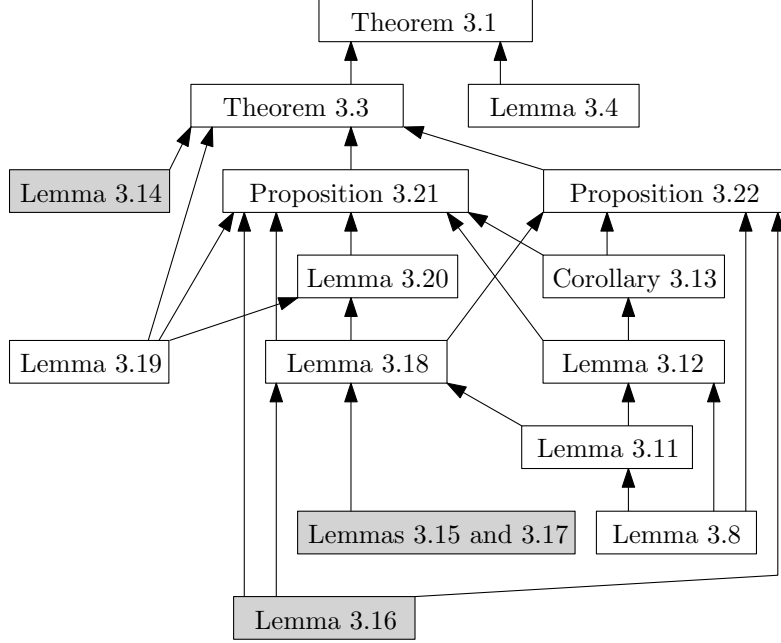


Figure 3.1: Flow summary. The shaded boxes correspond to computer-assisted results.

3.2 Proof of Theorem 3.1

We now apply Theorem 3.3 to obtain a superlinear lower bound on the number of 5-holes in a given set of n points. It clearly suffices to prove the statement for the case in which $n = 2^t$ for some integer $t \geq 5^5$.

We prove by induction on $t \geq 5^5$ that the number of 5-holes in an arbitrary set P of $n = 2^t$ points is at least $f(t) := c \cdot 2^t t^{4/5} = c \cdot n \log_2^{4/5} n$ for some absolute constant $c > 0$. For $t = 5^5$, we have $n > 10$ and, by the result of Harborth [Har78], there is at least one 5-hole in P . If the constant c is sufficiently small, then $f(t) = c \cdot n \log_2^{4/5} n \leq 1$ and we have at least $f(t)$ 5-holes in P , which constitutes our base case.

For the inductive step we assume that $t > 5^5$. We first partition P with a line ℓ into two sets A and B of size $n/2$ each. Then we further partition A and B into smaller sets using the following well-known lemma, which can be considered as a generalization of the well-known Ham-Sandwich cut and is for example implied by a result of Steiger and Zhao [SZ10, Theorem 1].

Lemma 3.4 (Generalized Ham-Sandwich Cut, [SZ10]). *Let $P' = A' \cup B'$ be an ℓ -divided set and let r be a positive integer such that $r \leq |A'|, |B'|$. Then there is a line that is disjoint from P' and that determines an open halfplane h with $|A' \cap h| = r = |B' \cap h|$.*

We set $r := \lfloor \log_2^{1/5} n \rfloor$, $s := \lfloor n/(2r) \rfloor$, and apply Lemma 3.4 iteratively in the following way to partition P into islands P_1, \dots, P_{s+1} of P so that for every $i \in \{1, \dots, s\}$, the sizes of $P_i \cap A$ and $P_i \cap B$ are exactly r . Let $P'_0 := P$. For every $i = 1, \dots, s$, we consider a line that is disjoint

from P'_{i-1} and that determines an open halfplane h with $|P'_{i-1} \cap A \cap h| = r = |P'_{i-1} \cap B \cap h|$. Such a line exists by Lemma 3.4 applied to the ℓ -divided set P'_{i-1} . We then set $P_i := P'_{i-1} \cap h$, $P'_i := P'_{i-1} \setminus P_i$, and continue with $i + 1$. Finally, we set $P_{s+1} := P'_s$. Figure 3.2 gives an illustration for the case $r = 7$.




Figure 3.2: An illustration of the proof of Theorem 3.1.

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

We thus further assume that for more than $s/2$ islands P_i , neither of the sets $P_i \cap A$ nor $P_i \cap B$ is in convex position. Since $r = \lfloor \log_2^{1/5} n \rfloor \geq 5$, Theorem 3.3 implies that there is an ℓ -divided 5-hole in each such P_i . Thus there is an ℓ -divided 5-hole in P_i for more than $s/2$ islands P_i . Since each P_i is an island of P and since $s = \lfloor n/(2r) \rfloor$, we have more than $s/2 \geq \lfloor n/(4r) \rfloor$ ℓ -divided 5-holes in P . As $|A| = |B| = n/2 = 2^{t-1}$, there are at least $f(t-1)$ 5-holes in A and at least $f(t-1)$ 5-holes in B by the inductive assumption. Since A and B are separated by the line ℓ , we have at least

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

5-holes in P . The right side of the above expression is at least $f(t) = cnt^{4/5}$ because the inequality $cn(t-1)^{4/5} + n/(4t^{1/5}) \geq cnt^{4/5}$ is equivalent to $(t-1)^{4/5}t^{1/5} + 1/(4c) \geq t$, which is true if the constant c is sufficiently small, as $(t-1)^{4/5}t^{1/5} \geq t-1$. This finishes the proof of Theorem 3.1.

3.3 Proof of Theorem 3.2

In this section we improve the lower bounds on the minimum number of 3-holes and 4-holes. To this end we use the notion of generated holes as introduced by García [Gar12].

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

We say that a point from P is *extremal* in P if it lies on the boundary of $\text{conv}(P)$. A point from P that is not extremal is *inner* in P . García [Gar12] proved the following relationships between $h_3(P)$ and $h_{3|5}(P)$ and between $h_4(P)$ and $h_{4|5}(P)$.

Theorem 3.5 ([Gar12]). *Let P be a set of n points and let $\gamma(P)$ be the number of extremal points of P . Then the following two equalities are satisfied:*

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

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

The proofs of both parts of Theorem 3.2 are carried out by induction on n similarly to the proof of Theorem 3.1. The base cases follow from the fact that each set P of $n \geq 10$ points contains at least one 5-hole in P and thus a generated 3-hole in P and a generated 4-hole in P . For the inductive step, let $P = A \cup B$ be an ℓ -divided set of n points with $|A|, |B| \geq \lfloor \frac{n}{2} \rfloor$, where n is a sufficiently large positive integer.

To show part (i), it suffices to prove $h_{3|5}(P) \geq \Omega(n \log^{2/3} n)$ as the statement then follows from Theorem 3.5. We use the recursive approach from the proof of Theorem 3.1, where we choose $r = \lfloor \log_2^{1/3} n \rfloor$. In each step of the recursion we either obtain $\lfloor \frac{n}{4r} \rfloor$ pairwise disjoint r -holes in P or $\lfloor \frac{n}{4r} \rfloor$ pairwise disjoint ℓ -divided 5-holes in P .

In the first case, each r -hole in P admits $\binom{r}{3}$ 3-holes in P and, by Theorem 3.5, it contains $\binom{r}{3} - r^2 + 5r - r - 4$ generated 3-holes in P . Thus, in total, we count at least $\frac{n}{4r} \binom{r}{3} - O(nr) \geq \Omega(n \log^{2/3} n)$ generated 3-holes in P .

In the second case, we have at least $\lfloor \frac{n}{4r} \rfloor$ ℓ -divided 5-holes in P . Without loss of generality, we can assume that at least $\frac{1}{2} \lfloor \frac{n}{4r} \rfloor \geq \lfloor \frac{n}{8r} \rfloor$ of those ℓ -divided 5-holes in P contain at least two points to the right of ℓ , as we otherwise continue with the horizontal reflection of P , which has ℓ as the axis of reflection. Note that if a 5-hole is reflected, then it generates another 3-hole, however, the number of generated 3-holes both in A and in B remains the same after the reflection by Theorem 3.5. Hence we can apply this transformation. Therefore we have at least $\lfloor \frac{n}{8r} \rfloor$ ℓ -divided generated 3-holes in P and, analogously as in the proof of Theorem 3.1, we obtain

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

This finishes the proof of part (i).

The proof of part (ii) is almost identical. We choose $r = \lfloor \log_2^{1/4} n \rfloor$ and use the facts that every r -hole in P contains $\binom{r}{4} - \frac{r^2}{2} + \frac{7r}{2} - r - 3$ generated 4-holes in P and that every ℓ -divided 5-hole in P generates two 4-holes in P , at least one of which is ℓ -divided. This finishes the proof of Theorem 3.2.

3.4 Preliminaries for the proof of Theorem 3.3

Before proceeding with the proof of Theorem 3.3, we first introduce some notation and definitions, and state some immediate observations.

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

Sectors of polygons For an integer $k \geq 3$, let \mathcal{P} be a convex polygon with vertices p_1, \dots, p_k traced counterclockwise in this order. We denote by $S(p_1, \dots, p_k)$ the open convex region to the left of each of the three lines $\overline{p_1 p_2}$, $\overline{p_1 p_k}$, and $\overline{p_{k-1} p_k}$. We call the region $S(p_1, \dots, p_k)$ a *sector* of \mathcal{P} . Note that every convex k -gon defines exactly k sectors. Figure 3.3(a) gives an illustration.




Figure 3.3: (a) An example of sectors. (b) An example of a^* -wedges with $t = |A| - 1$. (c) An example of a^* -wedges with $t < |A| - 1$.

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

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

Observation 3.6. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole in P . Then the following conditions are satisfied.*


- (i) *Every sector of an ℓ -divided 4-hole in P is empty of points of P .*
- (ii) *If S is a sector of a 4-hole in A and S is empty of points of A , then S is empty of points of B .*

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

- (i) Neither A nor B is in convex position.
- (ii) For every extremal point x of C , one of the sets $(C \setminus \{x\}) \cap A$ and $(C \setminus \{x\}) \cap B$ is in convex position.

Note that every ℓ -critical set $C = A \cup B$ contains at least four points in each of A and B . Figure 3.4 shows some examples of ℓ -critical sets. If $P = A \cup B$ is an ℓ -divided set with neither A nor B in convex position, then there exists an ℓ -critical island of P . This can be seen by iteratively removing extremal points so that none of the parts is in convex position after the removal.

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

Figure 3.4: Examples of ℓ -critical sets.

non-convex a -wedge otherwise. The indices of the a -wedges are considered modulo $|A| - 1$. In particular, $W_0^{(a)} = W_{|A|-1}^{(a)}$ and $W_{|A|}^{(a)} = W_1^{(a)}$.

If A is not in convex position, we denote the rightmost inner point of A as a^* and write $t := t^{(a^*)}$ and $W_k := W_k^{(a^*)}$ for $k = 1, \dots, |A| - 1$. Recall that a^* is unique, since all points have distinct x -coordinates. Figures 3.3(b) and 3.3(c) give an illustration.

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

Observation 3.7. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position. Then the points a_1, \dots, a_{t-1} lie to the right of a^* and the points $a_t, \dots, a_{|A|-1}$ lie to the left of a^* .*

3.5 Proof of Theorem 3.3

First, we give a high-level overview of the main ideas of the proof of Theorem 3.3. We proceed by contradiction and we suppose that there is no ℓ -divided 5-hole in a given ℓ -divided set $P = A \cup B$ with $|A|, |B| \geq 5$ and with neither A nor B in convex position. If $|A|, |B| = 5$, then the statement follows from the result of Harborth [Har78]. Thus we assume that $|A| \geq 6$ or $|B| \geq 6$. We reduce P to an island Q of P by iteratively removing points from the convex hull until one of the two parts $Q \cap A$ and $Q \cap B$ contains exactly five points or Q is ℓ -critical with $|Q \cap A|, |Q \cap B| \geq 6$. If $|Q \cap A| = 5$ and $|Q \cap B| \geq 6$ or vice versa, then we reduce Q to an island of Q with eleven points and, using a computer-aided result (Lemma 3.14), we show that there is an ℓ -divided 5-hole in that island and hence in P . If Q is ℓ -critical with $|Q \cap A|, |Q \cap B| \geq 6$, then we show that $|A \cap \partial \text{conv}(Q)|, |B \cap \partial \text{conv}(Q)| \leq 2$ and that, if $|A \cap \partial \text{conv}(Q)| = 2$, then a^* is the only inner point of $Q \cap A$ and similarly for B (Lemma 3.19). Without loss of generality, we assume that $|A \cap \partial \text{conv}(Q)| = 2$ and thus a^* is the only inner point of $Q \cap A$. Using this assumption, we prove that $|Q \cap B| < |Q \cap A|$ (Proposition 3.21). By exchanging the roles of $Q \cap A$ and $Q \cap B$, we obtain $|Q \cap A| \leq |Q \cap B|$ (Proposition 3.22), which gives a contradiction.

To prove that $|Q \cap B| < |Q \cap A|$, we use three results about the sizes of the parameters w_1, \dots, w_t for the ℓ -divided set Q , that is, about the numbers of points of $Q \cap B$ in the a^* -wedges W_1, \dots, W_t of Q . We show that if we have $w_i = 2 = w_j$ for some $1 \leq i < j \leq t$, then $w_k = 0$ for some k with $i < k < j$ (Lemma 3.12). Further, for any three or four consecutive a^* -wedges whose union is convex and contains at least four points of $Q \cap B$, each of those a^* -wedges contains at most two such points (Lemma 3.18). Finally, we show that $w_1, \dots, w_t \leq 3$ (Lemma 3.20). The proofs of Lemmas 3.18 and 3.20 rely on some results about small ℓ -divided sets with computer-aided proofs (Lemmas 3.15, 3.16, and 3.17). Altogether, this is sufficient to show that $|Q \cap B| < |Q \cap A|$.

We now start the proof of Theorem 3.3 by showing that if there is an ℓ -divided 5-hole in the intersection of P with a union of consecutive a^* -wedges, then there is an ℓ -divided 5-hole in P .

Lemma 3.8. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position. For integers i, j with $1 \leq i \leq j \leq t$, let $W := \bigcup_{k=i}^j W_k$ and $Q := P \cap W$. If there is an ℓ -divided 5-hole in Q , then there is an ℓ -divided 5-hole in P .*

Proof. If W is convex then Q is an island of P and the statement immediately follows. Hence we assume that W is not convex. The region W is bounded by the rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$ and all points of $P \setminus Q$ lie in the convex region $\mathbb{R}^2 \setminus W$; see Figure 3.5.




Figure 3.5: Illustration of the proof of Lemma 3.8. (a) The point a_j is to the right of a^* . (b) The point a_j is to the left of a^* . (c) The hole H properly intersects the ray $\overrightarrow{a^*a_j}$. The boundary of the convex hull of H is drawn red and the convex hull of H' is drawn blue.

Since W is non-convex and every a^* -wedge contained in W intersects ℓ , at least one of the points a_{i-1} and a_j lies to the left of a^* . Moreover, the points a_i, \dots, a_{j-1} are to the right of a^* by Observation 3.7. Without loss of generality, we assume that a_{i-1} is to the left of a^* , as otherwise we consider the vertical reflection of the whole point set P .

If a_j is to the left of a^* , then we let h be the closed halfplane determined by the vertical line through a^* such that a_{i-1} and a_j lie in h . Otherwise, if a_j is to the right of a^* , then we let h be the closed halfplane determined by the line $\overrightarrow{a^*a_j}$ such that a_{i-1} lies in h . In either case, $h \cap A \cap Q = \{a^*, a_{i-1}, a_j\}$.

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

Suppose for contradiction that H properly intersects both rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$. Then H contains the points a_{i-1}, a_j, x, y, z for some points $x, y, z \in Q$, where $a_{i-1}x$ intersects $\overrightarrow{a^*a_j}$, and a_jz intersects $\overrightarrow{a^*a_{i-1}}$. Observe that z is to the left of $\overrightarrow{a_{i-1}a^*}$ and that x is to the right of $\overrightarrow{a_ja^*}$. If a_j lies to the right of a^* , then z is to the left of a^* , and thus z is in A ; see Figure 3.5(a). However, this is impossible as z also lies in h . Hence, a_j lies to the left of a^* ; see Figure 3.5(b). As x and z are both to the right of a^* , the point a^* is inside the convex quadrilateral $\square(a_{i-1}, a_j, x, z)$. This contradicts the assumption that H is a 5-hole in Q .

So assume that H properly intersects exactly one of the rays $\overrightarrow{a^*a_{i-1}}$ and $\overrightarrow{a^*a_j}$, say $\overrightarrow{a^*a_j}$; see Figure 3.5(c). In this case, H contains a_{i-1} . The interior of the triangle $\triangle(a^*, a_{i-1}, a_j)$ is empty of points of Q , since the triangle is contained in h . Moreover, $\text{conv}(H)$ cannot intersect the line that determines h both strictly above and strictly below a^* . Thus, all remaining points of $H \setminus \{a_{i-1}\}$ lie to the right of $\overrightarrow{a_{i-1}a^*}$ and to the right of $\overrightarrow{a_ja^*}$. If H is empty of points of $P \setminus Q$,

we are done. Otherwise, we let $H' := (H \setminus \{a_{i-1}\}) \cup \{p'\}$ where $p' \in P \setminus Q$ is a point inside $\triangle(a^*, a_{i-1}, a_j)$ closest to $\overrightarrow{a_j a^*}$. Note that the point p' might not be unique. By construction, H' is an ℓ -divided 5-hole in P . An analogous argument shows that there is an ℓ -divided 5-hole in P if H properly intersects $\overrightarrow{a^* a_{i-1}}$.

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

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

In this subsection we consider an ℓ -divided set $P = A \cup B$ with A not in convex position. We consider the union W of consecutive a^* -wedges, each containing at most two points of B , and derive an upper bound on the number of points of B that lie in W if there is no ℓ -divided 5-hole in $P \cap W$; see Corollary 3.13.

Observation 3.9. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position. Let W_k be an a^* -wedge with $w_k \geq 1$ and $1 \leq k \leq t$ and let b be the leftmost point in $W_k \cap B$. Then the points a^* , a_{k-1} , b , and a_k form an ℓ -divided 4-hole in P .*

From Observation 3.6(i) and Observation 3.9 we obtain the following result.

Observation 3.10. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position and with no ℓ -divided 5-hole in P . Let W_k be an a^* -wedge with $w_k \geq 2$ and $1 \leq k \leq t$ and let b be the leftmost point in $W_k \cap B$. For every point b' in $(W_k \cap B) \setminus \{b\}$, the line $\overline{bb'}$ intersects the segment $a_{k-1}a_k$. Consequently, b is inside $\triangle(a_{k-1}, a_k, b')$, to the left of $\overline{a_k b'}$, and to the right of $\overline{a_{k-1} b'}$.*

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

Lemma 3.11. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position and with $|A|, |B| \geq 5$. Let W_i and W_{i+1} be consecutive a^* -wedges with $w_i = 2 = w_{i+1}$ and $1 \leq i < t$. Then there is an ℓ -divided 5-hole in P .*

Proof. The overall idea of the proof is as follows. We suppose for contradiction that there is no ℓ -divided 5-hole in P . Then we prove a sequence of structural facts on the layout of the points of P forced by this assumption. Eventually we show that the structure of the point set P resembles the point set from Figure 3.8((a)). In particular, we arrive at the conclusion that $|B| = 4$, which contradicts our assumption $|B| \geq 5$.

Suppose for contradiction that there is no ℓ -divided 5-hole in P . Let $W := W_i \cup W_{i+1}$ and let $Q := P \cap W$. By Lemma 3.8, there is also no ℓ -divided 5-hole in Q . We label the points in $B \cap W_i$ as b_{i-1} and b_i so that b_{i-1} is to the right of b_i . Similarly, we label the points in $B \cap W_{i+1}$ as b_{i+1} and b_{i+2} so that b_{i+2} is to the right of b_{i+1} . By Observation 3.10, the point a_i is to the right of $\overline{b_i b_{i-1}}$ and to the left of $\overline{b_{i+1} b_{i+2}}$. If the points $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are in convex position, then $a_i, b_{i+1}, b_{i+2}, b_{i-1}, b_i$ form an ℓ -divided 5-hole in P ; see Figure 3.6(a). Thus, we assume the points $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are not in convex position. Without loss of generality, we assume that the line $\overline{b_i b_{i-1}}$ intersects the segment $b_{i+1}b_{i+2}$, as otherwise we consider the horizontal reflection of the whole point set P .

We show that the segments $a_i b_{i-1}$ and $b_i b_{i+1}$ intersect. As $\overline{b_i b_{i-1}}$ intersects $a_i a_{i-1}$ and $b_{i+1} b_{i+2}$, the point b_{i-1} lies in the triangle $\triangle(b_i, b_{i+1}, b_{i+2})$. Moreover, b_{i-1} is to the right of $\overline{b_{i+1} b_i}$, a_i is to the left of $\overline{b_{i+1} b_i}$, b_i is to the left of $\overline{a_i b_{i-1}}$, and b_{i+1} is to the right of $\overline{a_i b_{i-1}}$. Consequently, the points $a_i, b_{i+1}, b_{i-1}, b_i$ form an ℓ -divided 4-hole in P , and, in particular, the segments $a_i b_{i-1}$ and $b_i b_{i+1}$ intersect, which we wanted to prove; see Figure 3.6(b).

The points $a_{i-1}, b_i, b_{i-1}, b_{i+2}$ are in convex position because a_{i-1} is the leftmost and b_{i+2} is the rightmost of those four points and because both a_{i-1} and b_{i+2} lie to the left of $\overline{b_i b_{i-1}}$.




Figure 3.6: (a) If $b_{i-1}, b_i, b_{i+1}, b_{i+2}$ are in convex position, then there is an ℓ -divided 5-hole in P . (b) The points $a^*, a_{i+1}, a_i, a_{i-1}$ form a 4-hole in P .

Moreover, the points $a_{i-1}, b_i, b_{i-1}, b_{i+2}$ form an ℓ -divided 4-hole in P as $\square(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ lies in W and $w_i = w_{i+1} = 2$.

We consider the four points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$. The point b_{i+2} is the rightmost of those four points. By Observation 3.10, b_{i+1} lies to the right of $\overline{a_i b_{i+2}}$ and a_{i+1} lies to the right of $\overline{b_{i+1} b_{i+2}}$. Since $b_{i-1} \in W_i$ and $b_{i+2} \in W_{i+1}$, the point b_{i-1} lies to the left of $\overline{a_i b_{i+2}}$. Thus, the clockwise order around b_{i+2} is $a_{i+1}, b_{i+1}, b_{i-1}$.

Suppose for contradiction that the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ form a convex quadrilateral. Due to the clockwise order around b_{i+2} , the convex quadrilateral is $\square(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$. The only points of P that can lie in the interior of this quadrilateral are a^*, a_{i-1}, a_i , and b_i . Since the triangle $\triangle(b_{i+2}, b_{i+1}, a_{i+1})$ is contained in W_{i+1} , it contains neither of the points a^*, a_{i-1}, a_i , and b_i . Since the triangle $\triangle(b_{i+2}, b_{i-1}, b_{i+1})$ is contained in the convex hull of B , it does not contain a^*, a_{i-1} , nor a_i . Moreover, as b_{i-1} lies in the triangle $\triangle(b_i, b_{i+1}, b_{i+2})$, the triangle $\triangle(b_{i+2}, b_{i-1}, b_{i+1})$ also does not contain b_i . Thus the quadrilateral $\square(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$ is empty of points of P . By Observation 3.6(i), the two sectors $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ and $S(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$ contain no point of P . Since every point of $B \setminus \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$ is either in $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ or in $S(b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1})$, we have $B = \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$. This contradicts the assumption that $|B| \geq 5$.

Therefore the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ are not in convex position. In particular, the point b_{i+1} lies in the triangle $\triangle(b_{i-1}, a_{i+1}, b_{i+2})$, since a_{i+1} is the leftmost and b_{i+2} is the rightmost of the points $b_{i+2}, b_{i-1}, b_{i+1}, a_{i+1}$ and since b_{i-1} lies in W_i . The red area in Figure 3.6(b) gives an illustration.

Consequently, the point a_{i+1} lies to the left of $\overline{b_{i+1} b_{i-1}}$. By Observation 3.6(i), the point a_{i+1} is not in the sector $S(b_{i+1}, b_{i-1}, b_i, a_i)$, as otherwise the points $b_{i+1}, b_{i-1}, b_i, a_i, a_{i+1}$ form an ℓ -divided 5-hole in P . Thus the point a_{i+1} lies to the left of $\overline{a_i b_i}$; see Figure 3.6(b).

The points $a^*, a_{i+1}, a_i, a_{i-1}$ do not form a 4-hole in P because otherwise b_i lies in the sector $S(a_{i-1}, a^*, a_{i+1}, a_i)$ and forms a 5-hole together with $a_{i-1}, a^*, a_{i+1}, a_i$, which is impossible by Observation 3.6(ii).

Therefore the points $a^*, a_{i+1}, a_i, a_{i-1}$ are not in convex position. Now we show that a^* is inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$. The point a_i is not inside $\triangle(a_{i-1}, a_{i+1}, a^*)$, since, by Observation 3.7, a_i is to the right of a^* and since a^* is the rightmost inner point of A . Since a_{i-1} is to the left of $\overline{a^* a_i}$ and a_{i+1} is to the right of $\overline{a^* a_i}$, a^* is the inner point of $a^*, a_{i+1}, a_i, a_{i-1}$. Figure 3.7 gives an illustration.

Since $|B| \geq 5$, there is another a^* -wedge besides W_i and W_{i+1} that intersects ℓ . Now we show that all points of $B \setminus Q$ lie in a^* -wedges below W_{i+1} . The rays $\overrightarrow{b_i a_{i-1}}$ and $\overrightarrow{b_{i-1} b_{i+2}}$ both start in W_i and then leave W_i . Moreover, the segment $b_i a_{i-1}$ intersects ℓ and $b_{i-1} b_{i+2}$ intersects $\overrightarrow{a^* a_i}$. As both b_i and b_{i-1} lie to the right of $\overline{a_{i-1} b_{i+2}}$, all points of $B \setminus Q$ that lie in an a^* -wedge above W_i also lie in the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$. We recall that, by Observation 3.6(i), the




Figure 3.7: Location of the points of $A \setminus Q$.

sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ is empty of points of P . Hence all points of $B \setminus Q$ lie in a^* -wedges below W_{i+1} .

Claim I. *We have $i = 1$. That is, W_i is the topmost a^* -wedge that intersects ℓ .*

By Observation 3.7, a_{i+1} lies to the right of a^* . Since a_i and a_{i+1} are both to the right of a^* and since a^* is inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$, the point a_{i-1} is to the left of a^* . By Observation 3.7, we have $i = 1$. This proves Claim I.

Claim II. *All points of $A \setminus Q$ lie to the left of $\overline{a_{i+1}a_i}$, to right of $\overline{a_{i+1}b_{i+1}}$, and to the right of $\overline{a^*a_{i+1}}$.*

The violet area in Figure 3.7 gives an illustration where the remaining points of $A \setminus Q$ lie. We recall that the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$ (red shaded area in Figure 3.7) is empty of points of P . By Observation 3.9, both sets $\{a^*, a_i, b_i, a_{i-1}\}$ and $\{a^*, a_{i+1}, b_{i+1}, a_i\}$ form ℓ -divided 4-holes in P . By Observation 3.6(i), the two sectors $S(a^*, a_i, b_i, a_{i-1})$ (green shaded area in Figure 3.7) and $S(a^*, a_{i+1}, b_{i+1}, a_i)$ (blue shaded area in Figure 3.7) are thus empty of points of P . Therefore, no point of $A \setminus Q$ lies to the left of $\overline{a_{i+1}b_{i+1}}$. Since W is non-convex, every point of P that is to the left of $\overline{a^*a_{i+1}}$ lies in Q . Thus every point of $A \setminus Q$ lies to the right of $\overline{a^*a_{i+1}}$. Moreover, no point a of $A \setminus Q$ lies to the right of $\overline{a_{i+1}a_i}$ (gray area in Figure 3.7) because otherwise, a_{i+1} is an inner point of $\triangle(a_i, a^*, a)$, which is impossible since a^* is the rightmost inner point of A and a_{i+1} is to the right of a^* . This finishes the proof of Claim II.

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

Claim III. *The points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ -divided 4-hole in P .*

We consider a_{i+2} and show that the points $a_{i+1}, a^*, a_{i-1}, a_{i+2}$ are in convex position. It suffices to show that a_{i+2} does not lie in the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ because of the cyclic order of $A \setminus \{a^*\}$ around a^* . Recall that a^* lies inside the triangle $\triangle(a_{i-1}, a_{i+1}, a_i)$, that b_{i+1} lies inside the triangle $\triangle(a_i, a_{i+1}, b_{i+2})$, and that b_{i-1} lies inside the triangle $\triangle(a_{i-1}, a_i, b_{i+2})$. Since the triangles $\triangle(a_{i-1}, a_{i+1}, a_i)$, $\triangle(a_i, a_{i+1}, b_{i+2})$, and $\triangle(a_{i-1}, a_i, b_{i+2})$ are oriented counterclockwise along the boundary, the point a_i lies inside $\triangle(a_{i-1}, a_{i+1}, b_{i+2})$. Thus also the points a^*, b_i, b_{i+1} lie in the triangle $\triangle(a_{i-1}, a_{i+1}, b_{i+2})$. Consequently, the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ is contained in the union of the sectors $S(a_{i+1}, b_{i+1}, a_i, a^*)$ (blue shaded area in Figure 3.7) and $S(a^*, a_i, b_i, a_{i-1})$ (green shaded area in Figure 3.7). Thus a_{i+2} does not lie in the triangle $\triangle(a^*, a_{i-1}, a_{i+1})$ and the points $a_{i+1}, a^*, a_{i-1}, a_{i+2}$ are in convex position.

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

By the choice of a'_{i-1} , the quadrilateral $\square(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$ is empty of points of P . Since a_{i+1} and a_{i+2} are consecutive in the order around a^* , no point of A lies in the sector $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$. By Observation 3.6(ii), the sector $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$ (gray shaded area in Figure 3.8(a)) is empty of points of P . Since the sector $S(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is a subset of $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$, the sector $S(a_{i+1}, a^*, a_{i-1}, a_{i+2})$ is empty of points of P .




Figure 3.8: (a) Location of the points of $B \setminus Q$. (b) The point a_{i+1} lies to the left of a_i .

We show that a_{i+1} is to the left of a_i and to the right of a_{i+2} . Recall that a_i lies to the right of a^* and to the left of b_i . The point b_i lies to the left of $\overline{a^*a_i}$ and the point a_{i+1} lies to the right of this line; see Figure 3.8(b). The point a_{i+1} then lies to the left of a_i , since we know already that a_{i+1} lies to the left of $\overline{a_i b_i}$. Recall that a_{i+1} is to the right of a^* . Consequently, the point a_{i+2} lies to the left of a_{i+1} , as a_{i+2} lies to the right of $\overline{a^*a_{i+1}}$ and to the left of $\overline{a_{i+1}a_i}$ by Claim II.

Now we are ready to prove that the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ -divided 4-hole in P (green area in Figure 3.8(a)). Recall that b_{i+2} and a_{i+2} both lie to the right of $\overline{a_{i+1}b_{i+1}}$, and that a_{i+2} is the leftmost and b_{i+2} is the rightmost of those four points. Altogether, we see that the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ are in convex position. The four sectors $S(b_{i+2}, a_{i-1}, b_i, b_{i-1})$ (red shaded area in Figure 3.8(a)), $S(b_{i-1}, b_i, a_i, b_{i+1})$ (orange shaded area in Figure 3.8(a)), $S(b_{i+1}, a_i, a^*, a_{i+1})$ (blue shaded area in Figure 3.8(a)), and $S(a_{i+1}, a^*, a'_{i-1}, a_{i+2})$ (gray shaded area in Figure 3.8(a)) contain the quadrilateral $\square(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$ (green area in Figure 3.8(a)). The sectors are empty of points of P by Observation 3.6(i). Consequently, the convex quadrilateral $\square(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$ is an ℓ -divided 4-hole in P . This concludes the proof of Claim III.

To finish the proof, recall that all points of $B \setminus Q$ lie in a^* -wedges below W_{i+1} as $i = 1$ by Claim I. Since a_{i+2} is to the left of a_{i+1} , the line $\overline{a_{i+2}a_{i+1}}$ intersects ℓ above $\ell \cap W_{i+2}$. The line $\overline{a_{i+1}b_{i+1}}$ also intersects ℓ above $\ell \cap W_{i+2}$, since a_{i+1} and b_{i+1} both lie in W_{i+1} . From $i = 1$, every point of $B \setminus Q$ is to the right of $\overline{a_{i+2}a_{i+1}}$ and to the right of $\overline{a_{i+1}b_{i+1}}$. Since the points $b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2}$ form an ℓ -divided 4-hole in P by Claim III, Observation 3.6(i) implies that the sector $S(b_{i+2}, b_{i+1}, a_{i+1}, a_{i+2})$ is empty of points of P . Thus every point of $B \setminus Q$ lies to the left of $\overline{b_{i+1}b_{i+2}}$. Since $\overline{b_{i+1}b_{i+2}}$ intersects $\ell \cap W_{i+1}$ above $\ell \cap a_{i+1}b_{i+1}$ and since b_{i-1} lies to the

left of b_{i+2} and to the left of $\overline{b_{i+1}b_{i+2}}$, every point of $B \setminus Q$ lies to the left of $\overline{b_{i-1}b_{i+2}}$ and to the right of b_{i+2} , and thus in the sector $S(a_{i-1}, b_i, b_{i-1}, b_{i+2})$. However, by Observation 3.6(i), this sector is empty of points of P . Thus we obtain $B = \{b_{i-1}, b_i, b_{i+1}, b_{i+2}\}$, which contradicts the assumption $|B| \geq 5$. This concludes the proof of Lemma 3.11. \square

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

Lemma 3.12. *Let $P = A \cup B$ be an ℓ -divided set with A not in convex position and with $|A| \geq 5$ and $|B| \geq 6$. Let W_i, \dots, W_j be consecutive a^* -wedges with $1 \leq i < j \leq t$, $w_i = 2 = w_j$, and $w_k = 1$ for every k with $i < k < j$. Then there is an ℓ -divided 5-hole in P .*

Proof. For $i = j - 1$, the statement follows by Lemma 3.11. Thus we assume $j \geq i + 2$. That is, we have at least three consecutive a^* -wedges. Suppose for contradiction that there is no ℓ -divided 5-hole in P . Let $W := \bigcup_{k=i}^j W_k$ and $Q := P \cap W$. By Lemma 3.8, there is also no ℓ -divided 5-hole in Q . Note that $|Q \cap B| = j - i + 3$. Also observe that $|Q \cap A| = j - i + 2$ if $a_{i-1} = a_j = a_t$ and $|Q \cap A| = j - i + 3$ otherwise. We label the points in $B \cap W_i$ as b_{i-1} and b_i so that b_{i-1} is to the right of b_i . Further, we label the unique point in $B \cap W_k$ as b_k for each $i < k < j$, and the two points in $B \cap W_j$ as b_j and b_{j+1} so that b_{j+1} is to the right of b_j ; see Figure 3.9.




Figure 3.9: An illustration of a^* -wedges W_i, \dots, W_j in the proof of Lemma 3.12.

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

The claim clearly holds for points from $B \cap W_k$. Thus it suffices to prove the claim only for points from $B \cap W_{k-1}$, as for points from $B \cap W_{k+1}$ it follows by symmetry. Since $i < k < j$, Observation 3.7 implies that the points a_{k-1} and a_k are both to the right of a^* .

We now distinguish the following two cases.

1. The point a_{k-2} is to the left of $\overline{a^* a_k}$; see Figure 3.10(a). Since a^* is the rightmost inner point of A , a_{k-1} does not lie inside $\triangle(a^*, a_k, a_{k-2})$ and thus $\square(a_{k-2}, a^*, a_k, a_{k-1})$ is a 4-hole in P . All points of $B \cap W_{k-1}$ lie to the right of $\overline{a^* a_{k-2}}$ and to the left of $\overline{a_{k-2} a_{k-1}}$. By Observation 3.6(ii), no point of $B \cap W_{k-1}$ lies in the sector $S(a_{k-2}, a^*, a_k, a_{k-1})$ (red shaded area in Figure 3.10(a)) and thus all points of $B \cap W_{k-1}$ are to the right of $\overline{a_k a_{k-1}}$.
2. The point a_{k-2} is to the right of $\overline{a^* a_k}$; see Figure 3.10(b). Since a_{k-1} and a_k are to the right of a^* and since a_{k-2} is to the left of $\overline{a^* a_{k-1}}$ and to the right of $\overline{a^* a_k}$, the point a_{k-2}




Figure 3.10: An illustration of the proof of Claim I.

is to the left of a^* . By Observation 3.7, we have $k = 2$. That is, W_{k-1} is the topmost a^* -wedge that intersects ℓ .

There is another a^* -wedge below W_{k+1} , since otherwise $|B| = |B \cap (W_{k-1} \cup W_k \cup W_{k+1})| \leq 2 + 1 + 2 = 5$, which is impossible according to the assumption $|B| \geq 6$. By Observation 3.7, the point a_{k+1} is to the right of a^* . Moreover, since a^* is the rightmost inner point of A , the point a_k does not lie inside the triangle $\triangle(a^*, a_{k+1}, a_{k-1})$. The points $a^*, a_{k+1}, a_k, a_{k-1}$ then form a 4-hole in P , which has a^* as the leftmost point.

By definition, all points of $B \cap W_{k-1}$ lie to the left of $\overrightarrow{a^* a_{k-1}}$. As the ray $\overrightarrow{a^* a_{k+1}}$ intersects ℓ , all points of $B \cap W_{k-1}$ lie also to the left of $\overrightarrow{a^* a_{k+1}}$. By Observation 3.6(ii), no point of $B \cap W_{k-1}$ lies in the sector $S(a^*, a_{k+1}, a_k, a_{k-1})$. Thus all points of $B \cap W_{k-1}$ lie to the right of $\overrightarrow{a_k a_{k-1}}$.

This finishes the proof of Claim I.

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

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

Due to Claim I, the points b_{i-1} and b_i are both to the right of $\overrightarrow{a_{i+1} a_i}$. Thus the points a_i and a_{i+1} are both extremal points of those four points. Also the point b_{i-1} is extremal, since it is the rightmost of those four points. The point b_i does not lie inside the triangle $\triangle(a_{i+1}, a_i, b_{i-1})$, since, by Observation 3.10, b_i lies to the left of $\overrightarrow{a_i b_{i-1}}$. To finish the proof of Claim II, it suffices to observe that the triples (b_{i-1}, b_i, a_i) , (b_{i-1}, b_i, a_{i+1}) , (b_{i-1}, a_i, a_{i+1}) , and (b_i, a_i, a_{i+1}) are all oriented counterclockwise.

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

Suppose for contradiction that b_{i+1} lies to the left of $\overrightarrow{b_i b_{i-1}}$. We consider the five points $a_{i-1}, a_i, b_{i-1}, b_i, b_{i+1}$; see Figure 3.11. By Claim I, the points b_{i-1}, b_i , and b_{i+1} lie to the right of $\overrightarrow{a_i a_{i-1}}$. Moreover, since b_{i-1} and b_i lie in W_i and since b_{i+1} lies in W_{i+1} , the points b_{i-1} and b_i both lie to the left of $\overrightarrow{a_i b_{i+1}}$. By Observation 3.10, the point a_{i-1} lies to the left of $\overrightarrow{b_i b_{i-1}}$ and b_{i+1} is to the right of b_{i-1} . Consequently, the points b_{i-1} and b_i lie in the triangle $\triangle(a_{i-1}, a_i, b_{i+1})$. Altogether, the points a_{i-1}, b_i, b_{i-1} , and b_{i+1} are in convex position.

By Claim I, the points b_{i-1} and b_{i+1} lie to the right of $\overrightarrow{a_{i+1} a_i}$. Moreover, since b_{i-1} is to the left of b_{i+1} and to the left of $\overrightarrow{a_i b_{i+1}}$, the points b_{i+1}, b_{i-1}, a_i , and a_{i+1} are in convex position. Since there are no further points in W_i and W_{i+1} , the sets $\{a_{i-1}, b_i, b_{i-1}, b_{i+1}\}$ and $\{b_{i+1}, b_{i-1}, a_i, a_{i+1}\}$ are ℓ -divided 4-holes in P . By Observation 3.6(i), the point b_{i+2} lies neither in $S(a_{i-1}, b_i, b_{i-1}, b_{i+1})$ nor in $S(b_{i+1}, b_{i-1}, a_i, a_{i+1})$. Recall that the ray $\overrightarrow{b_{i-1} b_{i+1}}$ intersects $\overrightarrow{a^* a_i}$ and the ray $\overrightarrow{b_i a_{i-1}}$ does not intersect $\overrightarrow{a^* a_i}$. Therefore b_{i+2} is to the right of $\overrightarrow{a_i a_{i+1}}$. This contradicts Claim I and finishes the proof of Claim III.

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




Figure 3.11: An illustration of the proof of Claim III.

Recall the labeling of the points in $B \cap W$; see Figure 3.9. We show by induction on k that

- (i) the points b_{i-1} , b_{k-1} , a_{k-1} , and a_k form a counterclockwise-oriented convex quadrilateral, which has b_{i-1} as the rightmost point, and
- (ii) the point b_k lies inside this convex quadrilateral and, in particular, to the left of $\overline{a_k b_{i-1}}$.

Claim IV then clearly follows.

For the base case, we consider $k = i + 1$. By Claim II, the points b_{i-1} , b_i , a_i , and a_{i+1} form a counterclockwise-oriented convex quadrilateral. By definition, b_{i-1} is the rightmost of those four points. Figure 3.12(a) gives an illustration. The point b_{i+1} lies to the right of $\overline{a_{i+1} a_i}$ and, by Claim III, to the right of $\overline{b_i b_{i-1}}$. Moreover, since b_{i+1} lies in W_{i+1} , it lies to the right of $\overline{a_i b_i}$. By Observation 3.6(i), b_{i+1} does not lie in the sector $S(b_{i-1}, b_i, a_i, a_{i+1})$. Consequently, b_{i+1} lies inside the quadrilateral $\square(b_{i-1}, b_i, a_i, a_{i+1})$.




Figure 3.12: (a) An illustration of the proof of Claim IV. (b) An illustration of the proof of Lemma 3.12.

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

Recall that the points b_{k-1} and b_k both lie to the right of $\overline{a_k a_{k-1}}$ and that b_{k-1} is to the left of $\overline{a_{k-1} b_{i-1}}$. Since $b_k \in W_k$, the point b_k lies to the right of $\overline{a_{k-1} b_{i-1}}$. Therefore the clockwise order of $\{b_{k-1}, b_{i-1}, b_k\}$ around a_{k-1} is b_{k-1}, b_{i-1}, b_k . Since b_{i-1} is not contained in $W_{k-1} \cup W_k$, the point b_{i-1} is not contained in the triangle $\triangle(a_{k-1}, b_k, b_{k-1})$. Consequently, the points a_{k-1} , b_k , b_{i-1} , b_{k-1} form a convex quadrilateral and, in particular, b_k lies to the right of $\overline{b_{k-1} b_{i-1}}$. Figure 3.12(a) gives an illustration. Since b_k lies in W_k , it lies to the right of $\overline{a_{k-1} b_{k-1}}$. By Observation 3.6(i), the point b_k does not lie in the sector $S(b_{i-1}, b_{k-1}, a_{k-1}, a_k)$. Thus b_k lies inside the quadrilateral $\square(b_{i-1}, b_{k-1}, a_{k-1}, a_k)$. This finishes the proof of Claim IV.

Using Claim IV, we now finish the proof of Lemma 3.12, by finding an ℓ -divided 5-hole in the island Q and thus obtaining a contradiction with the assumption that there is no ℓ -divided 5-hole in P . In the following, we assume, without loss of generality, that b_{j+1} is to the right of b_{i-1} . Otherwise we can consider a vertical reflection of P .

We consider the polygon \mathcal{P} through the points $b_{i-1}, b_{j-1}, a_{j-1}, b_j, b_{j+1}$ and we show that \mathcal{P} is convex and empty of points of Q . See Figure 3.12(b) for an illustration. This will give us an ℓ -divided 5-hole in Q .

We show that \mathcal{P} is convex by proving that every point of $\{b_{i-1}, b_{j-1}, a_{j-1}, b_j, b_{j+1}\}$ is a convex vertex of \mathcal{P} . The point a_{j-1} is a convex vertex of \mathcal{P} because it is the leftmost point in \mathcal{P} . The point b_{i-1} is a convex vertex of \mathcal{P} because all points of \mathcal{P} lie to the right of a^* and b_{i-1} is the topmost point in the clockwise order around a^* . The point b_{j+1} is a convex vertex of \mathcal{P} because b_{j+1} is the rightmost point in \mathcal{P} by Claim IV and by the assumption that b_{j+1} is to the right of b_{i-1} . The point b_{j-1} is a convex vertex of \mathcal{P} because b_{j-1} lies to the left of $\overline{a_{j-1}b_{j+1}}$ by Claim IV while b_j and b_{j+1} both lie to the right of this line. The point b_j is a convex vertex of \mathcal{P} because, by Observation 3.10, b_j lies to the right of $\overline{a_{j-1}b_{j+1}}$ while b_{j-1} and b_{i-1} both lie to the right of this line. Consequently, \mathcal{P} is a convex pentagon with vertices from both A and B . Moreover, by Claim IV, all points b_k with $i < k < j$ lie to the left of $\overline{a_k b_{i-1}}$. Since b_i is to the left of $\overline{b_{j-1} b_{i-1}}$, \mathcal{P} is thus empty of points of Q , which gives us a contradiction with the assumption that there is no ℓ -divided 5-hole in P . \square

We now use Lemma 3.12 to show the following upper bound on the total number of points of B in a sequence W_i, \dots, W_j of consecutive a^* -wedges with $w_i, \dots, w_j \leq 2$.

Corollary 3.13. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole, with A not in convex position, and with $|A| \geq 5$ and $|B| \geq 6$. For $1 \leq i \leq j \leq t$, let W_i, \dots, W_j be consecutive a^* -wedges with $w_k \leq 2$ for every k with $i \leq k \leq j$. Then $\sum_{k=i}^j w_k \leq j - i + 2$.*

Proof. Let n_0, n_1 , and n_2 be the number of a^* -wedges from W_i, \dots, W_j with 0, 1, and 2 points of B , respectively. Due to Lemma 3.12, we can assume that between any two a^* -wedges from W_i, \dots, W_j with two points of B each, there is an a^* -wedge with no point of B . Thus $n_2 \leq n_0 + 1$. Since $n_0 + n_1 + n_2 = j - i + 1$, we have $\sum_{k=i}^j w_k = 0n_0 + 1n_1 + 2n_2 = (j - i + 1) + (n_2 - n_0) \leq j - i + 2$. \square

3.5.2 Computer-assisted results

We now provide lemmas that are key ingredients in the proof of Theorem 3.3. All these lemmas have computer-aided proofs. Each result was verified by two independent implementations, which are also based on different abstractions of point sets; see below for details.

Lemma 3.14. *Let $P = A \cup B$ be an ℓ -divided set with $|A| = 5$, $|B| = 6$, and with A not in convex position. Then there is an ℓ -divided 5-hole in P .*

Lemma 3.15. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole in P , $|A| = 5$, $4 \leq |B| \leq 6$, and with A in convex position. Then for every point a of A , every convex a -wedge contains at most two points of B .*

Lemma 3.16. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole in P , $|A| = 6$, and $|B| = 5$. Then for each point a of A , every convex a -wedge contains at most two points of B .*

Lemma 3.17. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole in P , $5 \leq |A| \leq 6$, $|B| = 4$, and with A in convex position. Then for every point a of A , if the non-convex a -wedge is empty of points of B , every a -wedge contains at most two points of B .*

To prove these lemmas, we employ an exhaustive computer search through all combinatorially different sets of $|P| \leq 11$ points in the plane. Since none of these statements depends on the actual coordinates of the points but only on the relative positions of the points, we distinguish point sets only by orientations of triples of points (cf. Chapter 2).

We wrote two independent programs to verify Lemmas 3.14 to 3.17. Both programs are available online [Scha, Bal].

The first implementation is based on programs from my bachelor's theses [Sch13, Sch14]. For our verification purposes we reduced the framework from there to a very compact implementation [Scha]. The program uses the order type database [AAK02, AK06], which stores all order types realizable as point sets of size up to 11. As mentioned in Chapter 2.3, the statements can also be verified using the database of *abstract* order types. The running time of each of the programs in this implementation does not exceed two hours on a standard computer.

The second implementation [Bal] generates all possible abstract order types using a simple depth-first search algorithm and verifies the conditions from our lemmas. The running time of each of the programs in this implementation takes up to a few hundreds of hours.

3.5.3 Applications of the computer-assisted results

Here we present some applications of the computer-assisted results from Chapter 3.5.2.

Lemma 3.18. *Let $P = A \cup B$ be an ℓ -divided set with no ℓ -divided 5-hole in P , with $|A| \geq 6$, and with A not in convex position. Then the following two conditions are satisfied.*

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

Proof. To show part (i), let $W := W_i \cup W_{i+1} \cup W_{i+2}$, $A' := A \cap W$, $B' := B \cap W$, and $P' := A' \cup B'$. Since W is convex, P' is an island of P and thus there is no ℓ -divided 5-hole in P' . Note that $|A'| = 5$ and A' is in convex position. If $|B'| \leq 5$, then every convex a^* -wedge in P' contains at most two points of B' by Lemma 3.15 applied to P' . So assume that $|B'| \geq 6$. If necessary, we remove points from P' from the right to obtain $P'' = A' \cup B''$, where B'' contains exactly six points of B' . Note that there is no ℓ -divided 5-hole in P'' , since P'' is an island of P' . By Lemma 3.15, each a^* -wedge in P'' contains exactly two points of B'' . Let \tilde{B} be the set of points of B that are to the left of the rightmost point of B'' , including this point, and let $\tilde{P} := A \cup \tilde{B}$. Note that $B'' \subseteq \tilde{B}$. Since $|B''| = 6$ and since $W \cap \tilde{B} = B''$, each of the a^* -wedges W_i, W_{i+1}, W_{i+2} contains exactly two points of \tilde{B} . The a^* -wedges W_i, W_{i+1} , and W_{i+2} are also a^* -wedges in \tilde{P} . Thus, Lemma 3.11 applied to \tilde{P} and W_i, W_{i+1} then gives us an ℓ -divided 5-hole in \tilde{P} . From the choice of \tilde{P} , we then have an ℓ -divided 5-hole in P , a contradiction.

To show part (ii), let $W := W_i \cup W_{i+1} \cup W_{i+2} \cup W_{i+3}$, $A' := A \cap W$, $B' := B \cap W$, and $P' := A' \cup B'$. Since W is convex, P' is an island of P and thus there is no ℓ -divided 5-hole in P' . Note that $|A'| = 6$ and A' is in convex position. If $|B'| = 4$, then the statement follows from Lemma 3.17 applied to P' since a^* is an extremal point of P' . If $|B'| = 5$, then the statement follows from Lemma 3.16 applied to P' and thus we can assume $|B'| \geq 6$. Suppose for contradiction that $w_j \geq 3$ for some $i \leq j \leq i+3$. If necessary, we remove points from P from the right to obtain P'' so that $B'' := P'' \cap B$ contains exactly six points of $W \cap B$. By applying part (i) for P'' and $W_i \cup W_{i+1} \cup W_{i+2}$ and $W_{i+1} \cup W_{i+2} \cup W_{i+3}$, we obtain that $|B'' \cap W_i|, |B'' \cap W_{i+3}| = 3$ and $|B'' \cap W_{i+1}|, |B'' \cap W_{i+2}| = 0$. Let b be the rightmost point from $P'' \cap W$. By Lemma 3.16 applied to $W \cap (P'' \setminus \{b\})$, there are at most two points of $B'' \setminus \{b\}$ in every a^* -wedge in $W \cap (P'' \setminus \{b\})$. This contradicts the fact that either $|(B'' \cap W_i) \setminus \{b\}| = 3$ or $|(B'' \cap W_{i+3}) \setminus \{b\}| = 3$. \square

3.5.4 Extremal points of ℓ -critical sets

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

In this section, we consider an ℓ -critical set $C = A \cup B$ with $|A|, |B| \geq 5$. We first show that C has at most two extremal points in A and at most two extremal points in B . Later, under the assumption that there is no ℓ -divided 5-hole in C , we show that $|B| \leq |A| - 1$ if A contains two extremal points of C (Chapter 3.5.4) and that $|B| \leq |A|$ if B contains two extremal points of C (Chapter 3.5.4).

Lemma 3.19. *Let $C = A \cup B$ be an ℓ -critical set. Then the following statements are true.*

- (i) *If $|A| \geq 5$, then $|A \cap \partial \text{conv}(C)| \leq 2$.*
- (ii) *If $A \cap \partial \text{conv}(C) = \{a, a'\}$, then a^* is the only inner point in A and every point of $A \setminus \{a, a'\}$ lies in the convex region spanned by the lines $\overline{a^*a}$ and $\overline{a^*a'}$ that does not have any of a and a' on its boundary.*
- (iii) *If $A \cap \partial \text{conv}(C) = \{a, a'\}$, then the a^* -wedge that contains a and a' contains no point of B .*

By symmetry, analogous statements hold for B .

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

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

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

We remark that the assumption $|A| \geq 5$ in part (i) of Lemma 3.19 is necessary. In fact, arbitrarily large ℓ -critical sets with only four points in A and with three points of A on $\partial \text{conv}(C)$ exist, and analogously for B . Figure 3.4(c) gives an illustration.

Lemma 3.20. *Let $C = A \cup B$ be an ℓ -critical set with no ℓ -divided 5-hole in C and with $|A| \geq 6$. Then $w_i \leq 3$ for every $1 < i < t$. Moreover, if $|A \cap \partial \text{conv}(C)| = 2$, then $w_1, w_t \leq 3$.*

Proof. Recall that, since C is ℓ -critical, we have $|B| \geq 4$. Let i be an integer with $1 \leq i \leq t$. First, we assume that there is no point in $A \cap \partial \text{conv}(C)$, which lies outside of W_i . By Lemma 3.19(i), we have $|A \cap \partial \text{conv}(C)| \in \{1, 2\}$. If $|A \cap \partial \text{conv}(C)| = 1$, then, since there is no point in $A \cap \partial \text{conv}(C)$, we have $i \in \{1, t\}$ and there is nothing to prove for W_i . In the remaining case $|A \cap \partial \text{conv}(C)| = 2$, Lemma 3.19(iii) gives $W_i \cap B = \emptyset$, as there is no point in $A \cap \partial \text{conv}(C)$.

Hence there is a point a in $A \cap \partial \text{conv}(C)$, which lies outside of W_i . We consider $C' := C \setminus \{a\}$. Since C is an ℓ -critical set, $A' := C' \cap A$ is in convex position. Thus, there is a non-convex a^* -wedge W' of C' . Since W' is non-convex, all other a^* -wedges of C' are convex. Moreover, since W' is the union of the two a^* -wedges of C that contain a , all other a^* -wedges of C' are also a^* -wedges of C . Let W be the union of all a^* -wedges of C that are not contained in W' . Note that W is convex and contains at least $|A| - 3 \geq 3$ a^* -wedges of C . Since $|A| \geq 6$, the statement follows from Lemma 3.18(i). \square

Two extremal points of C in A

Proposition 3.21. *Let $C = A \cup B$ be an ℓ -critical set with no ℓ -divided 5-hole in C , with $|A|, |B| \geq 6$, and with $|A \cap \partial \text{conv}(C)| = 2$. Then $|B| \leq |A| - 1$.*

Proof. Since $|A \cap \partial \text{conv}(C)| = 2$, Lemma 3.20 implies that $w_i \leq 3$ for every $1 \leq i \leq t$. Let a and a' be the two points in $A \cap \partial \text{conv}(C)$. By Lemma 3.19(ii), all points of $A \setminus \{a, a'\}$ lie in the convex region R spanned by the lines $\overline{a^*a}$ and $\overline{a^*a'}$ that does not have any of a and a' on its boundary. That is, without loss of generality, $a = a_{h-1}$ and $a' = a_h$ for some $1 \leq h \leq |A| - 1$ and, by Lemma 3.19(iii), we have $w_h = 0$. Since all points of $A \setminus \{a, a'\}$ lie in the convex region R , the regions $W := \text{cl}(\mathbb{R}^2 \setminus (W_{h-1} \cup W_h))$ and $W' := \text{cl}(\mathbb{R}^2 \setminus (W_h \cup W_{h+1}))$ are convex. Here $\text{cl}(X)$ denotes the closure of a set $X \subseteq \mathbb{R}^2$. Recall that the indices of the a^* -wedges are considered modulo $|A| - 1$ and that \mathbb{R}^2 is the union of all a^* -wedges.

First, suppose for contradiction that $|A| = 6$. There are exactly five a^* -wedges W_1, \dots, W_5 , and only four of them can contain points of B , since $w_h = 0$. We can apply Lemma 3.18(i) to W and to W' . An easy case analysis shows that either $w_i \leq 2$ for every $1 \leq i \leq t$ or $w_{h-1}, w_{h+1} = 3$ and $w_i = 0$ for every $i \notin \{h-1, h+1\}$. In the first case, Corollary 3.13 implies that $|B| \leq 5$ and in the latter case Lemma 3.16 applied to $P \setminus \{b\}$, where b is the rightmost point of B , gives $|B| \leq 5$, a contradiction to $|B| \geq 6$. Hence, we assume $|A| \geq 7$.

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

There are $|A| - 1 \geq 6$ a^* -wedges and, in particular, W and W' are both unions of at least four a^* -wedges. For every W_i with $w_i = 3$ and $1 \leq i \leq t$, the a^* -wedge W_i is either contained in W or in W' . Thus we can find four consecutive a^* -wedges $W_k, W_{k+1}, W_{k+2}, W_{k+3}$ whose union is convex and contains W_i . Lemma 3.18(ii) implies that each of $W_k, W_{k+1}, W_{k+2}, W_{k+3}$ except of W_i is empty of points of B . This finishes the proof of Claim I.

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

Let $S := (w_1, \dots, w_j)$ and let S' be the subsequence of S obtained by removing every 1-entry from S . If S contains only 1-entries, the statement clearly follows. Thus we can assume that S' is non-empty. Recall that, by Lemma 3.20, S' contains only 0-, 2-, and 3-entries, since $w_i \leq 3$ for all $1 \leq i \leq t$. Due to Claim I, there are at least three consecutive 0-entries between every pair of nonzero entries of S' that contains a 3-entry. Together with Lemma 3.12, this implies that there is at least one 0-entry between every pair of 2-entries in S' .

By applying the following iterative procedure, we show that $\sum_{s \in S'} s \leq |S'| + 1$. While there are at least two nonzero entries in S' , we remove the first nonzero entry s from S' . If $s = 2$, then we also remove the 0-entry from S' that succeeds s in S . If $s = 3$, then we also remove the two consecutive 0-entries from S' that succeed s in S' . The procedure stops when there

is at most one nonzero element s' in the remaining subsequence S'' of S' . If $s' = 3$, then S'' contains at least one 0-entry and thus S'' contains at least $s' - 1$ elements. Since the number of removed elements equals the sum of the removed elements in every step of the procedure, we have $\sum_{s \in S'} s \leq |S'| + 1$. This implies

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

and finishes the proof of Claim II.

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

Two extremal points of C in B

Proposition 3.22. *Let $C = A \cup B$ be an ℓ -critical set with no ℓ -divided 5-hole in C , with $|A|, |B| \geq 6$, and with $|B \cap \partial \text{conv}(C)| = 2$. Then $|B| \leq |A|$.*

Proof. If $w_k \leq 2$ for all $1 \leq k \leq t$, then the statement follows from Corollary 3.13, since $|B| = \sum_{k=1}^t w_k \leq t + 1 \leq |A|$. Therefore we assume that there is an a^* -wedge W_i that contains at least three points of B . Let b_1, b_2 , and b_3 be the three leftmost points in $W_i \cap B$ from left to right. Without loss of generality, we assume that b_3 is to the left of $\overline{b_1 b_2}$. Otherwise we can consider a vertical reflection of P . Figure 3.13 gives an illustration.




Figure 3.13: An illustration of the proof of Proposition 3.22.

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

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

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

By Observation 3.10, both lines $\overline{b_1 b_2}$ and $\overline{b_1 b_3}$ intersect the segment $a_{i-1} a_i$. Since the segment $a_{i-1} b_1$ intersects ℓ and since b_1 is the leftmost point of $W_i \cap B$, all points of B' that lie to the left of $\overline{b_1 b_2}$ lie to the left of $\overline{a_{i-1} b_1}$. The four points a_{i-1}, b_1, b_2, b_3 form an ℓ -divided 4-hole in P , since a_{i-1} is the leftmost and b_3 is the rightmost point of a_{i-1}, b_1, b_2, b_3 and both a_{i-1} and b_3 lie to the left of $\overline{b_1 b_2}$. By Observation 3.6(i), the sector $S(a_{i-1}, b_1, b_2, b_3)$ is empty of points of P (green shaded area in Figure 3.13). Altogether, all points of B' that lie to the left of $\overline{b_1 b_2}$ are to the right of $\overline{b_2 b_3}$ and thus lie in R_1 .

Since the segment $a_i b_1$ intersects ℓ and since b_1 is the leftmost point of $W_i \cap B$, all points of B' that lie to the right of $\overline{b_1 b_2}$ lie to the right of $\overline{a_i b_1}$. By Observation 3.6(i), the sector $S(b_1, b_2, b_3, a_{i-1})$ is empty of points of P . Combining this with the fact that a^* is to the right of $\overline{a_{i-1} b_3}$, we see that a^* lies to the right of $\overline{b_1 b_2}$. Since b_1 and b_2 both lie to the left of $\overline{a^* a_i}$ and since a^* and a_i both lie to the right of $\overline{b_1 b_2}$, the points b_2, b_1, a^*, a_i form an ℓ -divided 4-hole in P .

By Observation 3.6(i), the sector $S(b_2, b_1, a^*, a_i)$ (blue shaded area in Figure 3.13) is empty of points of P . Altogether, all points of B' that lie to the right of $\overline{b_1 b_2}$ are to the right of $\overline{a^* a_i}$ and to the right of $\overline{a_i b_1}$ and thus lie in R_2 . This finishes the proof of Claim I.

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

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

We distinguish two cases. First, we assume that the points b_2, b_3, b_1, a_i are in convex position. Then b_2, b_3, b_1, a_i form an ℓ -divided 4-hole in P and, by Observation 3.6(i), the sector $S(b_2, b_3, b_1, a_i)$ is empty of points from P . Thus b_4 lies to the right of $\overline{b_2 b_3}$ and the statement follows.

Second, we assume that the points b_2, b_3, b_1, a_i are not in convex position. Due to Observation 3.10, b_2 and b_3 both lie to the right of $\overline{a_i b_1}$. Moreover, since b_3 is the rightmost of those four points, b_2 lies inside the triangle $\triangle(b_3, b_1, a_i)$. In particular, a_i lies to the right of $\overline{b_2 b_3}$. Therefore, since b_2 and b_3 are to the left of $\overline{a^* a_i}$, the line $\overline{b_2 b_3}$ intersects ℓ in a point p above $\ell \cap \overline{a^* a_i}$. Let q be the point $\ell \cap \overline{b_1 b_2}$. Note that q is to the left of $\overline{a^* a_i}$. The point b_4 is to the right of $\overline{b_2 b_3}$, as otherwise b_4 lies in $\triangle(p, q, b_2)$, which is impossible because the points p, q, b_2 are in W_i while b_4 is not. Altogether, b_2 is inside $\triangle(b_3, b_1, b_4)$ and this finishes the proof of Claim II.

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

By Observation 3.6(i), the sector $S(b_3, a_{i-1}, b_1, b_2)$ is empty of points of P and thus all points of $B' \cap R_1$ lie to the left of $\overline{a_{i-1} b_3}$ and, in particular, to the right of b_3 .

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

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

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

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

Suppose for contradiction that $|W \cap B| \geq 5$. Let $C' := C \cap W$. Note that $|C' \cap A| = 6$ and that a^*, a_{i-1}, a_i lie in C' . By Lemma 3.8, there is no ℓ -divided 5-hole in C' . We obtain C'' by removing points from C' from the right, if necessary, until $|C'' \cap B| = 5$. Since C'' is an island of C' , there is no ℓ -divided 5-hole in C'' . From Claim IV we know that b_1, b_2, b_3 are the three leftmost points in C and thus lie in C'' . We apply Lemma 3.16 to C'' and, since b_1, b_2, b_3 lie in a convex a^* -wedge of C'' , we obtain a contradiction. This finishes the proof of Claim V.

We now complete the proof of Proposition 3.22. First, we assume that $1 \leq i \leq 4$. Let $W := W_1 \cup W_2 \cup W_3 \cup W_4$. By Claim V, $|W \cap B| \leq 4$. Claim IV implies that $w_k \leq 2$ for every

k with $5 \leq k \leq t$. By Corollary 3.13, we have

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

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

Finally, we assume that $5 \leq i \leq t-4$. Let $W := W_{i-3} \cup W_{i-2} \cup W_{i-1} \cup W_i$. Note that W is convex, since $2 \leq i-3$ and $i < t$. By Lemma 3.18(ii), we have $w_{i-3} + w_{i-2} + w_{i-1} + w_i \leq 3$ and $w_i + w_{i+1} + w_{i+2} + w_{i+3} \leq 3$. By Claim IV, $w_k \leq 2$ for all k with $1 \leq k \leq i-4$. Thus, by Corollary 3.13, $\sum_{k=1}^{i-4} w_k \leq i-3$. Similarly, we have $\sum_{k=i+4}^t w_k \leq t-i-2$. Altogether, we obtain that

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

□

3.5.5 Finalizing the proof of Theorem 3.3

We are now ready to prove Theorem 3.3. Namely, we show that for every ℓ -divided set $P = A \cup B$ with $|A|, |B| \geq 5$ and with neither A nor B in convex position there is an ℓ -divided 5-hole in P .

Suppose for the sake of contradiction that there is no ℓ -divided 5-hole in P . By the result of Harborth [Har78], every set P of ten points contains a 5-hole in P . In the case $|A|, |B| = 5$, the statement then follows from the assumption that neither of A and B is in convex position.

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

- (i) One of the parts $Q \cap A$ and $Q \cap B$ has only five points.
- (ii) Q is an ℓ -critical island of P with $|Q \cap A|, |Q \cap B| \geq 6$.

In case (i), we have $|Q \cap A| = 5$ or $|Q \cap B| = 5$. We can assume by symmetry that $|Q \cap A| = 5$ and $|Q \cap B| \geq 6$. We let Q' be the union of $Q \cap A$ with the six leftmost points of $Q \cap B$. Since $Q \cap A$ is not in convex position, Lemma 3.14 implies that there is an ℓ -divided 5-hole in Q' , which is also an ℓ -divided 5-hole in Q , since Q' is an island of Q . However, this is impossible as then there is an ℓ -divided 5-hole in P because Q is an island of P .

In case (ii), we have $|Q \cap A|, |Q \cap B| \geq 6$. There is no ℓ -divided 5-hole in Q , since Q is an island of P . By Lemma 3.19(i), we can assume without loss of generality that $|A \cap \partial \text{conv}(Q)| = 2$, as $|A \cap \partial \text{conv}(Q)| + |B \cap \partial \text{conv}(Q)| \geq 3$ and thus $|A \cap \partial \text{conv}(Q)|$ and $|B \cap \partial \text{conv}(Q)|$ cannot be both smaller than 2. Then it follows from Proposition 3.21 that $|Q \cap B| < |Q \cap A|$. By exchanging the roles of $Q \cap A$ and $Q \cap B$ and by applying Proposition 3.22, we obtain that $|Q \cap A| \leq |Q \cap B|$, a contradiction. This finishes the proof of Theorem 3.3.

3.6 Necessity of Assumptions

In this section, we present some point configurations which show that the statements of our computer-assisted lemmas are best possible in some sense. In fact, finding the final setup of the lemmas was an iterative process and we luckily had a complete proof in the end.

3.6.1 Necessity of the assumptions in Theorem 3.3

In the statement of Theorem 3.3 we require that the ℓ -divided set $P = A \cup B$ satisfies $|A|, |B| \geq 5$. We now show that those requirements are necessary in order to guarantee an ℓ -divided 5-hole in P by constructing an arbitrarily large ℓ -critical set $C = A \cup B$ with $|A| = 4$ and with no ℓ -divided 5-hole in C .

Proposition 3.23. *For every integer $n \geq 5$, there exists an ℓ -critical set $C = A \cup B$ with $|A| = 4$, $|B| = n$, and with no ℓ -divided 5-hole in C .*

Proof. First, we consider the case where n is odd. Let $p^+ = (0, 1)$ and $p^- = (0, -1)$ be two auxiliary points and let $\ell^+ = \{(x, y) \in \mathbb{R}^2 : y = x/4\}$ and $\ell^- = \{(x, y) \in \mathbb{R}^2 : y = -x/4\}$ be two auxiliary lines. We place the point $b'_1 = (2, -1/2)$ on the line ℓ^- and the auxiliary point $q = (2, 1/2)$ on the line ℓ^+ . For $i = 2, \dots, n$, we iteratively let b'_i be the intersection of the line ℓ^+ with the segment $p^+b'_{i-1}$ if i is even and the intersection of ℓ^- with $p^-b'_{i-1}$ if i is odd. We place two points a_1 and a_2 sufficiently close to p^+ so that a_1 is above a_2 , the segment a_1a_2 is vertical with the midpoint p^+ , and all non-collinear triples (b'_i, b'_j, p^+) have the same orientation as (b'_i, b'_j, a_1) and (b'_i, b'_j, a_2) . Similarly, we place two points a_3 and a_4 sufficiently close to p^- so that a_3 is to the left of a_4 , the segment a_3a_4 lies on the line $\overline{p^-q}$ and has p^- as its midpoint, the point a_4 is to the left of b'_n , and all non-collinear triples (b'_i, b'_j, p^-) have the same orientation as (b'_i, b'_j, a_3) and (b'_i, b'_j, a_4) . Figure 3.14 gives an illustration.




Figure 3.14: The set C constructed in the proof of Proposition 3.23 for n odd.

We let A , B' , and B'_3 be the sets $\{a_1, a_2, a_3, a_4\}$, $\{b'_1, \dots, b'_n\}$, and $B' \setminus \{b'_3\}$, respectively. Note that the line $\overline{a_3a_4}$ intersects the segment $b'_1b'_3$. Since $\max_{a \in A} x(a) < \min_{b' \in B'} x(b')$, the sets A and B' are separated by a vertical line ℓ .

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

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

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

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

We now distinguish the following two cases.

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

If $H \cap A = A^-$, then the hole H contains no point from B^+ . This is because if there is a point b_i in $H \cap B^+$, then the point b_{i+1} lies in the interior of $\text{conv}(H)$. The point b_1 cannot lie in H because otherwise H is not in convex position as the line $\overline{a_3 a_4}$ separates b_1 from $B \setminus \{b_1\}$. Additionally, H contains at most two points from B^- , since otherwise H is not in convex position. Thus H contains at most four points of C , which is impossible.

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

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

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

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

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

3.6.2 Necessity of the assumptions in Lemmas 3.14 to 3.17

We remark that all the assumptions in the statements of Lemmas 3.14 to 3.17 are necessary; Figure 3.15(a) shows that the conditions $|B| = 5$ in Lemma 3.16 and the convexity of A in Lemma 3.17 are both necessary. The horizontal reflection of Figure 3.15(a) also shows the necessity of the assumption $|A| = 5$ in Lemma 3.14. It follows from the example in Figure 3.15(b) that the condition $|B| = 4$ cannot be omitted in Lemma 3.17, since there is an a -wedge with three points of B . The same point set without the point a' shows that the assumption $|B| \geq 4$ in Lemma 3.15 is necessary. The example from Figure 3.15(c) shows that the conditions $|B| = 6$ in Lemma 3.14, the convex position of A in Lemma 3.15, and $|A| = 6$ in Lemma 3.16 are all necessary. The same set without the point a shows that $|A| = 5$ in Lemma 3.15 is also needed and, if we remove the points a and a' , then the resulting point set shows that we need $5 \leq |A|$ in Lemma 3.17. We can make statements only about convex a -wedges in Lemmas 3.15 and 3.16, as there are counterexamples for the corresponding statements without the convexity condition. It suffices to consider so-called *double-chains*, which are point sets obtained by placing n points on each of the two branches of a hyperbola. Double-chains also show, that A cannot be in convex position in Lemma 3.14, and, that the non-convex a -wedge must be empty of points in B in Lemma 3.17.




Figure 3.15: Examples of points sets that witness tightness of Lemmas 3.14 to 3.17. All k -holes in these sets with $k \geq 5$ are highlighted in gray.

Chapter 4

Finding Holes using SAT Solvers

In this chapter we develop a framework for Boolean formulas in conjunctive normal form (CNF) to investigate various combinatorial properties of point sets using SAT solvers. As our main result, we obtain that every set of 17 points in general position admits two disjoint 5-holes.

Theorem 4.1 (Computer-assisted). *Every set of 17 points contains two disjoint 5-holes, hence $h(5, 5) = 17$.*

The proof of the upper bound is based on a SAT model which we describe in Chapter 4.4. The lower bound $h(5, 5) \geq 17$ is witnessed by the set of 16 points with no two disjoint 5-holes (taken from Hosono and Urabe [HU08]), which is depicted Figure 4.1¹.




Figure 4.1: A set of 16 points with no two disjoint 5-holes. This point set and the one by Hosono and Urabe [HU08, Figure 3] are of the same order type.

The computations for verifying Theorem 4.1 take about two hours on a single 3 GHz CPU using a modern SAT solver such as glucose (version 4.0)² or picosat (version 965)³. Moreover,

¹For readability, only particular edges are drawn in the figure; we refer the interested reader to [ABH⁺19b].

²<http://www.labri.fr/perso/lsimon/glucose/>, see also [AS09]

³<http://fmv.jku.at/picosat/>, see also [Bie08]

we have verified the output of glucose and picosat with the proof checking tool drat-trim⁴ (see Chapter 4.4.2).

In Chapter 4.1, we summarize the current state of the art for three-parametric values $h(k_1, k_2, k_3)$ and we present some new results that were obtained by utilizing the value $h(5, 5) = 17$. Moreover, we describe some direct consequences for multi-parametric values $h(k_1, \dots, k_l)$ in Chapter 4.2.

The basic idea behind our computer-assisted proofs is to encode point sets and disjoint holes only using triple orientations (see Chapter 2 and Chapter 4.3), and then to use a SAT solver to disprove the existence of sets with certain properties (see Chapter 4.4).

In the Chapter 4.5 we outline how our SAT model can be adapted to tackle related questions on point sets. For interior-disjoint holes, we show that every set of 15 points contains two interior-disjoint 5-holes. Also it is remarkable, that our SAT model can be used to prove $g(6) = 17$ with significantly smaller computation time than the original program from Szekeres and Peters [SP06]. Last but not least, we also outline how SAT solvers can be used to count occurrences of certain substructures (such as k -holes in point sets).

4.1 Three Disjoint Holes

For three parameters, most values $h(k_1, k_2, k_3)$ for $k_1, k_2, k_3 \leq 4$ and also the values $h(2, 3, 5) = 11$ and $h(3, 3, 5) = 12$ are known [HU08, YW15]. Tables 4.1 and 4.2 summarize the currently best known bounds for three-parametric values.

	2	3	4
2	8	9	11
3		10	12
4			14

Table 4.1: Values of $h(k_1, k_2, 4)$.

	2	3	4	5
2	10	11	11..14	17*
3		12	13..14	17..19*
4			15..17	17..23*
5				22*..27*

Table 4.2: Bounds for $h(k_1, k_2, 5)$.

The values $h(2, 2, 4)$, $h(3, 3, 4)$, and $h(2, 4, 4)$ have not been explicitly stated in literature. However, the former two can be derived directly from other values (cf. Table 1.2) as follows:

$$\begin{aligned} 8 &= 2 + 2 + 4 \leq h(2, 2, 4) \leq 2 + h(2, 4) = 8 \\ 10 &= 3 + 3 + 4 \leq h(3, 3, 4) \leq 3 + h(3, 4) = 10. \end{aligned}$$

To determine the value $h(2, 4, 4) = 11$, observe that $h(2, 4, 4) \leq 2 + h(4, 4) = 11$ clearly holds. Equality is witnessed by the *double circle* with 10 points (cf. Figure 4.2). This statement can be verified by computer or as follows: First, observe that no 4-hole contains two consecutive extremal points, thus every 4-hole contains at most two exterior points. Now consider two disjoint 4-holes. Since not both 4-holes can contain two extremal points, one of them contains two exterior points while the other one contains one exterior point. As illustrated in Figure 4.2, this configuration is unique up to symmetry and does not allow any further disjoint 2-hole. This completes the argument.

Also we could not find the value $h(2, 2, 5)$ in literature, however, using a SAT instance similar to the one for Theorem 4.1 one can also easily verify that $h(2, 2, 5) \leq 10$, and equality follows from $h(5) = 10$ [Har78]. One can also use the database of all (abstract) order types on 10 points, to verify the existence of those particular disjoint holes for all possible configurations of 10 points.

⁴<http://cs.utexas.edu/~marijn/drat-trim>, see also [WHH14]




Figure 4.2: The double circle on 10 points witnesses $h(2, 4, 4) > 10$.

We now use Theorem 4.1 to derive new bounds on the value $h(k, 5, 5)$ for $k = 2, 3, 4, 5$.

Corollary 4.2. *We have*

$$h(2, 5, 5) = 17, \quad 17 \leq h(3, 5, 5) \leq 19, \quad 17 \leq h(4, 5, 5) \leq 23, \quad \text{and} \quad 22 \leq h(5, 5, 5) \leq 27.$$

Proof. To show $h(2, 5, 5) \leq 17$, observe that, due to Theorem 4.1, every set of 17 points contains two disjoint 5 holes that are separated by a line ℓ . By the pigeonhole principle there are at least 9 points on one of the two sides of such a separating line ℓ . Again, using a SAT instance similar to the one for Theorem 4.1, one can easily verify that every set of 9 points with at least one 5-hole contains a 5-hole and a 2-hole that are disjoint. This completes the argument. We remark that one can also use the order type database of 9 points to verify this statement.

To show $h(3, 5, 5) \leq 2 \cdot h(3, 5) - 1 = 19$, observe that, due to Theorem 4.1, every set of 19 points contains two disjoint 5 holes that are separated by a line ℓ . Now there are at least 10 points on one side of such a separating line ℓ , and since $h(3, 5) = 10$, there is a 3-hole and a 5-hole that are disjoint on that particular side.

An analogous argument shows $h(4, 5, 5) \leq 2 \cdot h(4, 5) - 1 = 23$.

The set of 21 points depicted in Figure 4.3 witnesses $h(5, 5, 5) > 21$ (can be easily verified by computer), while $h(5, 5, 5) \leq h(5) + h(5, 5) = 27$. We remark that this point set was found using local search techniques, implemented in our framework *pyotlib*⁵. \square

4.2 Many Disjoint Holes

As introduced by Hosono and Urabe [HU01, HU08], we use the following notation: Given positive integers k and n , let $F_k(n)$ denote the maximum number of pairwise disjoint k -holes that can be found in every set of n points, that is,

$$F_k(n) := \max(\{0\} \cup \{t \in \mathbb{N} : h(k; t) \leq n\}) \quad \text{with} \quad h(k; t) := h(\overbrace{k, k, \dots, k}^{t \text{ parameters}}).$$

⁵The “**python order type library**” was initiated during the Bachelor’s studies of the author [Sch14] and provides many features to work with (abstract) order types such as local search techniques, realization or proving non-realizability of abstract order types, coordinate minimization and “beautification” for nicer visualizations.




Figure 4.3: A set of 21 points with no three disjoint 5-holes.

In the following, we revise and further improve results by Hosono and Urabe [HU01, HU08] and by Bárány and Károlyi [BK01]. The currently best bounds are the following:

$$\begin{aligned}
 F_k(n) &= \lfloor n/k \rfloor & \text{for } k = 1, 2, 3 \\
 3n/13 + o(n) &\leq F_4(n) < n/4 \\
 2n/17 - O(1) &\leq F_5(n) < n/6 \\
 n/h(6) - O(1) &\leq F_6(n) < n/12 \\
 F_k(n) &= 0 & \text{for } k \geq 7.
 \end{aligned}$$

Hosono and Urabe [HU01] showed that $F_4(n) \geq (3n - 1)/13$ holds for an infinite sequence of integers n . Moreover, since we have

$$h(k; s + t) \leq h(k; s) + h(k; t),$$

Fekete's subadditivity lemma (see for example [Sch03, Chapter 14.5]) asserts

$$\lim_{t \rightarrow \infty} \frac{h(k; t)}{t} = \inf_{t \in \mathbb{N}} \frac{h(k; t)}{t},$$

and consequently $3n/13 + o(n) \leq F_4(n)$ holds.

Concerning the lower bound on $F_5(n)$, Theorem 4.1 clearly implies that $F_5(n) \geq \lfloor 2n/17 \rfloor$ holds.

Concerning the upper bounds, it was remarked in [BK01] that $F_5(n) < n/6$ is not too difficult to prove but no explicit construction was given. We now outline how the upper bounds $F_5(n) < n/6$ and $F_6(n) < n/12$ can be obtained from the double circle on $2n$ points with an additional “center point”: Every 5-hole (6-hole) in this “dotted double circle” is incident to at

most 2 extremal points, and therefore, at most $2/3$ ($2/4$) of the exterior points – that is less than $5/6$ ($3/4$) of all points – can be covered by disjoint 5-holes (6-holes). An analogous statement shows that the dotted double circle on $4k+1$ points has no k disjoint 4-holes, hence $F_4(n) < n/4$. In particular, we obtain that $h(4, 4, 4, 4) = 18$ since $17 < h(4, 4, 4, 4) \leq 2h(4, 4) = 18$.

It is also worth to note that the double circle (sometimes with the additional center point, sometimes without) is a maximal configuration also for other settings; see for example [HU01, Figure 5] and [HU08, Figures 1(a), 4, 10(a), and 10(b)].

4.3 Encoding with Triple Orientations

As discussed in Chapter 2, point sets and disjoint holes can be encoded only using triple orientations. This combinatorial description allows us to get rid of the actual point coordinates and to only consider a discrete parameter-space. This is essential for our SAT model of the problem.

4.3.1 Signotope Axioms

The chirotope axioms can be checked in $\Theta(n^6)$ time, and, when a Boolean satisfiability instance (SAT) is modelled, $\Theta(n^6)$ constraints are needed. Moreover, it is sufficient to only check 5-tuples (cf. Theorem 3.6.2. [BLW⁺99]), which gives a slight improvement to $\Theta(n^5)$. We now describe an axiomatization which only requires $\Theta(n^4)$ conditions to be checked – the so-called signotope axioms.

Since we only consider point sets in general position, we can assume without loss of generality that in any set $S = \{s_1, \dots, s_n\}$ the points s_1, \dots, s_n have strictly increasing x -coordinates. Due to Felsner and Weil [FW01] (see also [BFK15]), for every 4-tuple s_i, s_j, s_k, s_l with $i < j < k < l$ the sequence

$$\chi_{ijk}, \chi_{ijl}, \chi_{ikl}, \chi_{jkl}$$

(index-triples are in lexicographic order) changes its sign at most once. These conditions are the *signotope axioms*. It is worth mentioning that these necessary conditions were also used in the computer-assisted proof for $g(6) = 17$ by Szekeres and Peters [SP06] and also later by Balko and Valtr [BV17], who refuted a natural strengthening of the Erdős–Szekeres conjecture introduced by Szekeres and Peters.

4.3.2 Increasing Coordinates and Cyclic Order

In the following, we see why we can assume, without loss of generality, that in every point set $S = \{s_1, \dots, s_n\}$ the following three conditions hold:

- the points s_1, \dots, s_n have increasing x -coordinates,
- in particular, s_1 is an extremal point, and
- the points s_2, \dots, s_n are sorted around s_1 .

When modeling a computer program, one can use these constraints (which do not affect the output of the program) to restrict the search space and to possibly get a speedup. This idea, however, is not new and was already used for the generation of the order type database [Kra03, AAK02, AK06].

Lemma 4.3. *Let $S = \{s_1, \dots, s_n\}$ be a point set where s_1 is extremal and s_2, \dots, s_n are sorted around s_1 . Then there is a point set $\tilde{S} = \{\tilde{s}_1, \dots, \tilde{s}_n\}$ of the same order type as S (in particular, $\tilde{s}_2, \dots, \tilde{s}_n$ are sorted around \tilde{s}_1) such that the points $\tilde{s}_1, \dots, \tilde{s}_n$ have increasing x -coordinates.*

Proof. We can apply an appropriate affine-linear transformation to S so that $s_1 = (0, 0)$ and $x_i, y_i > 0$ holds for $i \geq 2$. Moreover, we have that x_i/y_i is increasing for $i \geq 2$ since s_2, \dots, s_n are sorted around s_1 . Since S is in general position, there is an $\varepsilon > 0$ such that S and $S' := \{(0, \varepsilon)\} \cup \{s_2, \dots, s_n\}$ are of the same order type. We apply the projective transformation $(x, y) \mapsto (x/y, -1/y)$ to S' to obtain \tilde{S} . By the multilinearity of the determinant, we obtain

$$\det \begin{pmatrix} 1 & 1 & 1 \\ x_i & x_j & x_k \\ y_i & y_j & y_k \end{pmatrix} = y_i \cdot y_j \cdot y_k \cdot \det \begin{pmatrix} 1 & 1 & 1 \\ x_i/y_i & x_j/y_j & x_k/y_k \\ -1/y_i & -1/y_j & -1/y_k \end{pmatrix}.$$

Since the points in S' have positive y -coordinates, S' and \tilde{S} have the same triple orientations. Moreover, as $\tilde{x}_i = x'_i/y'_i$ is increasing for $i \geq 1$, the set \tilde{S} fulfills all desired properties. \square

It is worth to mention that the transformation $(x, y) \mapsto (x/y, -1/y)$ is the concatenation of the (inverse of the) *unit paraboloid duality transformation* and *unit circle duality transformation* which – under the given conditions – both preserve the triple orientations (see e.g. [Kra03, Chapters 1.3 and 2.2]).

4.4 SAT Model

In this section we describe the SAT model that we use to prove Theorem 4.1. The basic idea of the proof is to assume towards a contradiction that a point set $S = \{s_1, \dots, s_{17}\}$ with no two disjoint 5-holes exists. We formulate a SAT instance, where Boolean variables indicate whether triples are positively or negatively oriented and clauses encode the necessary conditions introduced in Chapters 2 and 4.3. Using a SAT solver we verify that the SAT instance has no solution and conclude that the point set S does not exist. This contradiction then completes the proof of Theorem 4.1.

It is folklore that satisfiability is NP-hard in general, thus it is challenging for SAT solvers to terminate in reasonable time for certain SAT instances. We now highlight the two crucial parts of our SAT model, which are indeed necessary for reasonable computation times: First, due to Lemma 4.3, we can assume without loss of generality that the points are sorted from left to right and also around the first point s_1 . Second, we teach the solver that every set of 10 consecutive points gives a 5-hole, that is, $h(5) = 10$ [Har78]. By dropping either of these two constraints (which only give additional information to the solver and do not affect the solution space), none of the tested SAT solvers terminated within days.

In the following, we give a detailed description of our SAT model. For the sake of readability, we refer to points also by their indices. Moreover, we use the relation “ $a < b$ ” simultaneously to indicate a larger index, a larger x -coordinate, and the later occurrence in the cyclic order around s_1 .

4.4.1 A Detailed Description

(1) Alternating axioms For every triple (a, b, c) , we introduce the variable $O_{a,b,c}$ to indicate whether the triple (a, b, c) is positively oriented. Since we have that

$$\chi_{a,b,c} = \chi_{b,c,a} = \chi_{c,a,b} = -\chi_{b,a,c} = -\chi_{a,c,b} = -\chi_{c,b,a},$$

we formulate clauses to assert

$$O_{a,b,c} = O_{b,c,a} = O_{c,a,b} \neq O_{b,a,c} = O_{a,c,b} = O_{c,b,a}$$

by using the fact $A = B \iff (\neg A \vee B) \wedge (A \vee \neg B)$, and $A \neq B \iff (A \vee B) \wedge (\neg A \vee \neg B)$.

(2) Signotope Axioms As described in Chapter 2.2, for every 4-tuple $a < b < c < d$, the sequence

$$\chi_{abc}, \chi_{abd}, \chi_{acd}, \chi_{bcd}$$

changes its sign at most once. Formally, to forbid such sign-patterns (that is, “ $- + -$ ” and “ $+ - +$ ”), we add the constraints

$$O_I \vee \neg O_J \vee O_K \quad \text{and} \quad \neg O_I \vee O_J \vee \neg O_K$$

for every lexicographically ordered triple of index triples, that is, $\{I, J, K\} \subset \binom{\{a,b,c,d\}}{3}$ with $I \prec J \prec K$.

(3) Sorted around first point Since the points are sorted from left to right and also around s_1 , we have that all triples $(1, a, b)$ are positively oriented for $1 < a < b$.

(4) Bounding segments For a 4-tuple a, b, c, d , we introduce the auxiliary variable $E_{a,b;c,d}$ to indicate whether the segment ab spanned by a and b bounds the convex hull $\text{conv}(\{a, b, c, d\})$. Since the segment ab bounds the convex hull $\text{conv}(\{a, b, c, d\})$ if and only if c and d lie on the same side of the line \overline{ab} , we add the constraints

$$\begin{aligned} \neg E_{a,b;c,d} &\vee O_{abc} \vee \neg O_{abd}, \\ \neg E_{a,b;c,d} &\vee \neg O_{abc} \vee O_{abd}, \\ E_{a,b;c,d} &\vee O_{abc} \vee O_{abd}, \\ E_{a,b;c,d} &\vee \neg O_{abc} \vee \neg O_{abd}. \end{aligned}$$

(5) 4-Gons and containments For every 4-tuple $a < b < c < d$, we introduce the auxiliary variable $G_{a,b,c,d}^4$ to indicate whether the points $\{a, b, c, d\}$ form a 4-gon. Moreover we introduce the auxiliary variable $I_{i;a,b,c}$ for every 4-tuple a, b, c, i with $a < b < c$ and $a < i < c$ to indicate whether the point i lies inside the triangular convex hull $\text{conv}(\{a, b, c\})$.

Four points $a < b < c < d$, sorted from left to right, form a 4-gon if and only if both segments ab and cd bound the convex hull $\text{conv}(\{a, b, c, d\})$. Moreover, if $\{a, b, c, d\}$ does not form a 4-gon, then either b lie inside the triangular convex hull $\text{conv}(\{a, c, d\})$ or c lies inside $\text{conv}(\{a, b, d\})$. Pause to note that a and d are the left- and rightmost points, respectively, and that not both points b and c can lie in the interior of $\text{conv}(\{a, b, c, d\})$. Formally, we assert

$$\begin{aligned} G_{a,b,c,d}^4 &= E_{a,b;c,d} \wedge E_{c,d;a,b}, \\ I_{b;a,c,d} &= \neg E_{a,b;c,d} \wedge E_{c,d;a,b}, \\ I_{c;a,b,d} &= E_{a,b;c,d} \wedge \neg E_{c,d;a,b}. \end{aligned}$$

(6) 3-Holes For every triple of points $a < b < c$, we introduce the auxiliary variable $H_{a,b,c}^3$ to indicate whether the points $\{a, b, c\}$ form a 3-hole. Since three points $a < b < c$ form a 3-hole if and only if every other point i lies outside the triangular convex hull $\text{conv}(\{a, b, c\})$, we add the constraint

$$H_{a,b,c}^3 = \bigwedge_{i \in S \setminus \{a,b,c\}} \neg I_{i;a,b,c}.$$

(7) 5-Holes For every 5-tuple $X = \{a, b, c, d, e\}$ with $a < b < c < d < e$, we introduce the auxiliary variable H_X^5 to indicate that the points from X form a 5-hole. It is easy to see that the points from X form a 5-hole if and only if every 4-tuple $Y \in \binom{X}{4}$ forms a 4-gon and if every triple $Y \in \binom{X}{3}$ forms a 3-hole. Therefore, we add the constraint

$$H_X^5 = \left(\bigwedge_{Y \in \binom{X}{4}} G_Y^4 \right) \wedge \left(\bigwedge_{Y \in \binom{X}{3}} H_Y^3 \right).$$

(8) Forbid disjoint 5-holes If there were two disjoint 5-holes X_1 and X_2 in our point set S , then – as discussed in Chapter 2 – we could find two points $a \in X_1$ and $b \in X_2$ such that the line \overline{ab} separates $X_1 \setminus \{a\}$ and $X_2 \setminus \{b\}$ – and this is what we have to forbid in our SAT model. Hence, for every pair of two points a, b we introduce the variables

- $L_{a,b}$ to indicate that there exists a 5-hole X containing the point a that lies to the left of the directed line \overrightarrow{ab} , that is, the triple (a, b, x) is positively oriented for every $x \in X \setminus \{a\}$, and
- $R_{a,b}$ to indicate that there exists a 5-hole X containing the point b that lies to the right of the directed line \overrightarrow{ab} , that is, the triple (a, b, x) is negatively oriented for every $x \in X \setminus \{b\}$.

For every 5-tuple X with $a \in X$ and $b \notin X$ we assert

$$L_{a,b} \vee \neg H_X \vee \left(\bigvee_{c \in X \setminus \{a\}} \neg O_{a,b,c} \right),$$

and for every 5-tuple X with $a \notin X$ and $b \in X$ we assert

$$R_{a,b} \vee \neg H_X \vee \left(\bigvee_{c \in X \setminus \{b\}} O_{a,b,c} \right).$$

Now we forbid that there are 5-holes on both sides of the line \overline{ab} by asserting

$$\neg L_{a,b} \vee \neg R_{a,b}.$$

(9) Harborth's result Harborth [Har78] has shown that every set of 10 points gives a 5-hole, that is, $h(5) = 10$. Consequently, there is a 5-hole X_1 in the set $\{1, \dots, 10\}$, and if $X_1 \subset \{1, \dots, 7\}$, then there is another 5-hole X_2 in the set $\{8, \dots, 17\}$. Analogously, if there is a 5-hole $X_3 \subset \{11, \dots, 17\}$, then there is another 5-hole X_4 in the set $\{1, \dots, 10\}$. Therefore, we can teach the SAT solver that

- there is a 5-hole X with $X \subset \{1, \dots, 10\}$,
- there is no 5-hole X with $X \subset \{1, \dots, 7\}$,
- there is a 5-hole X with $X \subset \{8, \dots, 17\}$, and
- there is no 5-hole X with $X \subset \{11, \dots, 17\}$.

We remark that the so-obtained SAT instance has $\Theta(n^5)$ variables and $\Theta(n^6)$ clauses. The source code of our python program which creates the instance is available online on our supplemental website [Schb].

4.4.2 Unsatisfiability and Verification

Having the satisfiability instance generated, we used the following command to create an unsatisfiability certificate:

```
glucose instance.cnf -certified -certified-output=proof.out
```

The certificate created by glucose was then verified using the proof checking tool drat-trim by the following command:

```
drat-trim instance.cnf proof.out
```

The execution of each of the two commands (glucose and drat-trim), took about 2 hours and the certificate used about 3.1 GB of disk space.

We have also used pycosat to prove unsatisfiability:

```
picosat instance.cnf -R proof.out
```

This command ran for about 6 hours and created a certificate of size about 2.1 GB. The verification of the certificate⁶ using drat-trim took about 9 hours.

4.5 Further Applications of the SAT Model

Interior-disjoint Holes: By slightly adapting the SAT model from Chapter 4.4, we managed to show that every set of 15 points contains two interior-disjoint 5-holes. Moreover, this bound is best possible because, for example, the set of 14 points depicted in Figure 4.4 contains no two interior-disjoint 5-holes. This further improves Theorem 3 from [HU18], which asserts that every set of 18 points contains two interior-disjoint 5-holes.




Figure 4.4: A set of 14 points with no two interior-disjoint 5-holes.

To be more specific on the changes of the SAT model for this variant: we slightly relaxed the constraints “(8) Forbid disjoint 5-holes” so that each of the two points a and b , which span a separating line ℓ , can be contained in holes from both sides. The program creating the SAT instance is also available on our website [Schb].

We remark that, analogously to Chapter 4.1, one could further improve the bounds for three interior-disjoint holes.

Classical Erdős–Szekeres: The computation time for the computer assisted proof by Szekeres and Peters [SP06] for $g(6) = 17$ was about 1500 hours. By slightly adapting the model from Chapter 4.4 we have been able to confirm $g(6) = 17$ using glucose and drat-trim with about one hour of computation time. To be more specific with the adaption of the model from Chapter 4.4:

- The constraints “(6) 3-Holes” are removed.

⁶In our experiments, pycosat wrote a comment “%RUPD32 ...” as first line in the RUP file. This line had to be removed manually to make the file parsable for drat-trim.

- The constraints “(7) 5-Holes” are adapted to “(7*) 6-Gons” simply by testing 6-tuples instead of 5-tuples and by dropping the requirement that “triples form 3-holes”.
- The constraints “(8) Forbid disjoint 5-holes” are removed.

Also this program is available on our website [Schb].

For determining the exact value of $g(7)$, however, further ideas or more advanced SAT solvers seem to be required.

Counting 5-Holes: It is also possible to count occurrences of certain substructures using SAT solvers. For example to find point sets with as few 5-holes as possible, we have introduced variables $X_{abcde;k}$ indicating whether the indices $1 \leq a < b < c < d < e \leq n$ form the k -th 5-hole in lexicographic order. In particular, using SAT solvers we have been able to show that every set of 16 points contains at least 11 5-holes. Our obtained results were presented in Table 1.1 of Chapter 1.

Part II

Arrangements of Pseudocircles

Chapter 5

Arrangements of Pseudocircles

The study of arrangements of pseudolines – which generalize arrangements of lines in a natural way¹ – was initiated in 1926 with an article of Levi [Lev26] where he proved the ‘Extension Lemma’ and studied triangular cells in arrangements. Since then arrangements of pseudolines were intensively studied and the handbook article on the topic [FG18] lists more than 100 references.

Arrangements of pseudocircles generalize arrangements of circles in the same vein as arrangements of pseudolines generalize arrangements of lines. To the best of our knowledge the study of arrangements of pseudocircles was initiated by Grünbaum [Grü80] in the 1970s. By stating a large number of conjectures he was hoping to attract the attention of researchers for the topic. The success of this program was limited and several of Grünbaum’s 45 year old conjectures remain unsettled.

In this thesis we give the first thorough study of circularizability – the problem of deciding which arrangements of pseudocircles are isomorphic to an arrangement of circles – and report on some progress regarding conjectures involving numbers of triangles and digons in arrangements of pseudocircles. Some of our results and new conjectures are based on a program that enumerates all arrangements of up to 7 pairwise intersecting pseudocircles.

A *pseudocircle* is a simple closed curve in the plane or on the sphere. An *arrangement of pseudocircles* is a collection of pseudocircles with the property that the intersection of any two of the pseudocircles is either empty or consists of two points where the curves cross. The left-hand side of Figure 5.1 gives an illustration. It is worth to mention that other authors also allow touching pseudocircles (see e.g. [ANP⁺04]).

The (*primal*) *graph of an arrangement* \mathcal{A} of pseudocircles has the intersection points of pseudocircles as *vertices*, the vertices split each of the pseudocircles into arcs, these are the *edges* of the graph. Note that this graph may have multiple edges and loop edges without vertices. The graph of an arrangement of pseudocircles comes with a plane embedding, the faces of this embedding are the *cells* of the arrangement.

In an arrangement \mathcal{A} of pseudocircles, we denote a cell with k crossings on its boundary as a k -*cell* and let $p_k(\mathcal{A})$ be the number of k -cells of \mathcal{A} . Following Grünbaum we call 2-cells *digons* and remark that some other authors call them *lenses*. 3-cells are *triangles*, 4-cells are *quadrangles*, and 5-cells are *pentagons*; see the right-hand side of Figure 5.1 for an illustration.

An arrangement \mathcal{A} of pseudocircles is

simple, if no three pseudocircles of \mathcal{A} intersect in a common point;

connected, if the graph of the arrangement is connected;

intersecting, if any two pseudocircles of \mathcal{A} intersect;

¹A *pseudoline* is a simple closed non-contractible curve in the projective plane. An *arrangement of pseudolines* is a collection of pseudolines such that any two intersect in exactly one point where they cross.




Figure 5.1: An arrangement of 5 pseudocircles (left) and with its k -cells highlighted (right). The arrangement has $p_2 = 6$ digons (blue), $p_3 = 4$ triangles (green), $p_4 = 8$ quadrangles (red), $p_5 = 0$ pentagons, and $p_6 = 4$ hexagons (black).

cylindrical, if there are two cells of the arrangement \mathcal{A} which are separated by each of the pseudocircles;

digon-free, if there is no cell of the arrangement which is incident to only two pseudocircles.

Note that every intersecting arrangement is connected. In this thesis we assume that arrangements of pseudocircles are simple unless explicitly stated otherwise.

Two arrangements \mathcal{A} and \mathcal{B} are *isomorphic* if they induce homeomorphic cell decompositions of the compactified plane, i.e., on the sphere. Stereographic projections can be used to map between arrangements of pseudocircles in the plane and arrangements of pseudocircles on the sphere. Figure 5.2² gives an illustration. Such projections are also considered isomorphisms. In particular, the isomorphism class of an arrangement of pseudocircles in the plane is closed under changes of the unbounded cell. In many cases, in particular in all our figures, arrangements of pseudocircles are embedded in the Euclidean plane, i.e., there is a distinguished outer/unbounded cell. An advantage of such a representation is that we can refer to the inner and outer side of a pseudocircle. Note that for every cylindrical arrangement of pseudocircles it is possible to choose the unbounded cell such that the intersection of the inner discs of all pseudocircles is non-empty.




Figure 5.2: A stereographic projection of an arrangement of two circles from the sphere to the plane.

²This 3D-figure was created using GeoGebra [H⁺19]

Figure 5.3 shows the three connected arrangements of three pseudocircles and Figure 5.4 shows the 21 connected arrangements of four pseudocircles. We call the unique digon-free intersecting arrangement of three (pseudo)circles the *Krupp*³. The second intersecting arrangement is the *NonKrupp*; this arrangement has digons. The non-intersecting arrangement is the *3-Chain*.




Figure 5.3: The 3 connected arrangements of 3 pseudocircles. (a) *Krupp*, (b) *NonKrupp*, (c) *3-Chain*.

Every triple of an arrangement of great-circles on the sphere induces a Krupp, hence, we call an arrangement of pseudocircles an *arrangement of great-pseudocircles* if every subarrangement induced by three pseudocircles is a Krupp. Some authors think of arrangements of great-pseudocircles when they speak about arrangements of pseudocircles, this is e.g. common practice in the theory of oriented matroids. In fact, arrangements of great-pseudocircles serve to represent rank 3 oriented matroids (cf. [BLW⁺99]). Planar partial cubes can be characterized as the duals of so-called ‘non-separating’ arrangements of pseudocircles, these are certain arrangements such that no triple forms a NonKrupp [AK16].

5.1 Circularizability

As introduced by Grünbaum (cf. [Grü80, page 68]), we call an arrangement of pseudocircles *circularizable* if there is an isomorphic arrangement of circles. In his book Grünbaum wrote that deciding circularizability appears to be a “very difficult” problem and, in fact, preceeding our work there have been only few results about circularizability of arrangements of pseudocircles.

Edelsbrunner and Ramos [ER97] presented an intersecting arrangement of 6 pseudocircles (with digons) which has no realization with circles, i.e., it is not circularizable; see the left-hand side of Figure 5.5. Linhart and Ortner [LO05] found a non-circularizable non-intersecting arrangement of 5 pseudocircles with digons; see the right-hand side of Figure 5.5. Linhart and Ortner also proved that every intersecting arrangement of at most 4 pseudocircles is circularizable, and Kang and Müller [KM14] extended the result by showing that all arrangements with at most 4 pseudocircles are circularizable. Kang and Müller also proved that deciding circularizability for connected arrangements is NP-hard. Actually, by reducing the problem of deciding stretchability of an arrangement of n pseudolines to a problem of deciding circularizability of an connected arrangement of $2n$ pseudocircles, they even showed ETR-completeness⁴. The complexity class ETR (“existential theory of the reals”) consists of problems, that can be reduced in polynomial time to solvability of a system of polynomial inequalities in several variables over the reals, and lies inbetween NP and PSPACE. The inclusion $\text{NP} \subseteq \text{ETR}$ clearly follows from the fact, that every Boolean formula can be written as a polynomial, and $\text{ETR} \subseteq \text{PSPACE}$ was shown by Canny [Can88]. Further background on ETR can be found in [Mat14, SŠ17] and in the Wikipedia article [Wikc].

In Chapter 7 we present the results of the first thorough study of circularizability. We show that there are exactly four non-circularizable arrangements of 5 pseudocircles (one of them was known before). In the set of 2131 digon-free intersecting arrangements of 6 pseudocircles we

³This name refers to the logo of the Krupp AG, a German steel company. Krupp was the largest company in Europe at the beginning of the 20th century.

⁴Figure 33 from [KM14] nicely illustrates the proof idea.

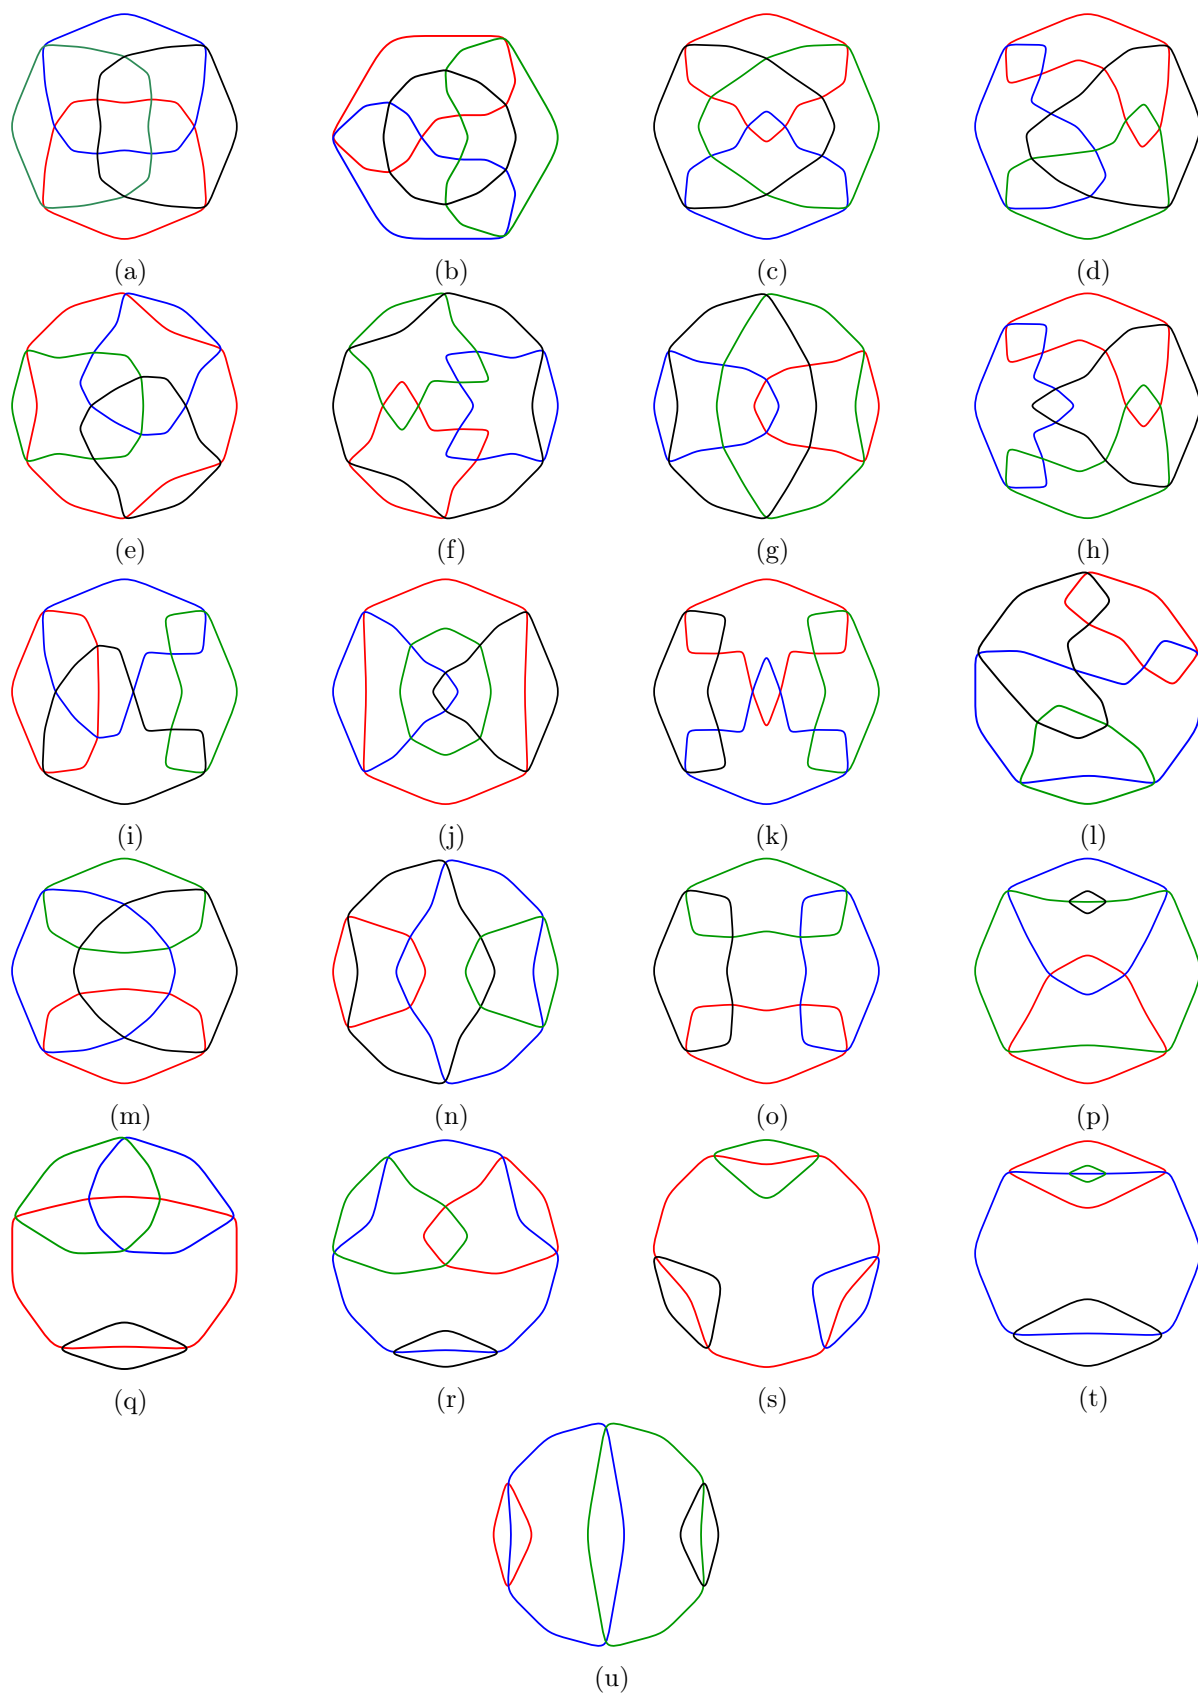


Figure 5.4: The 21 connected arrangements of 4 pseudocircles. The 8 first arrangements (a)–(h) are intersecting. The arrangements (a), (b), and (m) are digon-free. The arrangement (s) is the unique non-cylindrical.




Figure 5.5: The Edelsbrunner–Ramos arrangement (left) and the Linhart–Ortner arrangement (right).

identify the three non-circularizable examples. We also show non-circularizability of 8 additional arrangements of 6 pseudocircles which have a group of symmetries of size at least 4.

Most of our non-circularizability proofs depend on incidence theorems like Miquel’s. In other cases we contradict circularizability by considering a continuous deformation where the circles of an assumed circle representation grow or shrink in a controlled way.

The claims that we have all non-circularizable arrangements with the given properties are based on a program that generated all arrangements up to a certain size. Given the complete lists of arrangements, we used heuristics to find circle representations. Examples where the heuristics failed were examined by hand.

Furthermore, we will intensively study arrangements of great-pseudocircles – a very specific class of pseudocircles – and show that such an arrangement is circularizable if and only if it is realizable as an arrangement of great-circles. Based on this result, we then strengthen the hardness-result of Kang and Müller [KM14] in the following two ways:

- We reduce stretchability of an arrangement of n pseudolines to circularizability of an arrangement of n pseudocircles.
- The resulting arrangement is an arrangement of great-pseudocircles – a very specific class of pseudocircles⁵.

It is also worth mentioning that our hardness-reduction is somewhat more natural since it is based on an argument of sweeping planes and does not require any gadgets.

5.2 Triangular Cells

Conjecture 3.7 from Grünbaum’s monograph [Grü80] is: *Every (not necessarily simple) digon-free arrangement of n pairwise intersecting pseudocircles has at least $2n - 4$ triangles*. Grünbaum also provides examples of arrangements with $n \geq 6$ pseudocircles and $2n - 4$ triangles; see Figure 5.6.

Snoeyink and Hershberger [SH91] showed that the sweeping technique, which serves as an important tool for the study of arrangements of lines and pseudolines, can be adapted to work also in the case of arrangements of pseudocircles (cf. Section 6.3). They used sweeps to show that, in an intersecting arrangement, every pseudocircle is incident to two cells which are digons or triangles on both sides (two in the interior and two in the exterior). Therefore, $2p_2 + 3p_3 \geq 4n$ which implies that every intersecting digon-free arrangement of n pseudocircles has at least $4n/3$ triangles.

⁵Arrangements of great-pseudocircles are the most restrictive class considered in this thesis.




Figure 5.6: Illustration of the construction by Grünbaum [Grü80, Figure 3.28]: a digon-free arrangement of $n = 8$ pseudocircles with $p_3 = 12$ triangles.

Felsner and Kriegel [FK99] observed that the bound from [SH91] also applies to non-simple intersecting digon-free arrangements and gave examples of arrangements showing that the bound is tight on this class for infinitely many values of n . These examples disprove Grünbaum's conjecture in the non-simple case.

In Chapter 8 we present examples to disprove Grünbaum's conjecture in the general case. With a recursive construction based on an example with 12 pseudocircles and 16 triangles we obtain a family of intersecting digon-free arrangements with $p_3(\mathcal{A})/n \rightarrow 16/11 = 1.\overline{45}$. We expect that the lower bound $p_3(\mathcal{A}) \geq 4n/3$ is tight for infinitely many simple arrangements. It may however be true that all digon-free arrangements of n pairwise intersecting circles have at least $2n - 4$ triangles.

Furthermore, for pairwise intersecting arrangements with digons we have a lower bound of $p_3 \geq 2n/3$, and conjecture that $p_3 \geq n - 1$. Concerning the maximum number of triangles in pairwise intersecting arrangements of pseudocircles, we show that $p_3 \leq \frac{4}{3}\binom{n}{2} + O(n)$. This is essentially best possible because there are families of pairwise intersecting arrangements of n pseudocircles with $p_3 = \frac{4}{3}\binom{n}{2}$.

Chapter 6

Preliminaries

Stereographic projections map circles to circles (if we consider a line to be a circle containing the point at infinity), therefore, circularizability on the sphere and in the plane is the same concept. Arrangements of circles can be mapped to isomorphic arrangements of circles via Möbius transformations. In this context, the sphere is identified with the extended complex plane $\mathbb{C} \cup \{\infty\}$. Note that, for $n \geq 2$, the isomorphism class of an arrangement of n circles is not covered by Möbius transformations. Indeed, if \mathcal{C} is a simple arrangement of circles, then ε -perturbations of the circles in size and position will result in an isomorphic arrangement when ε is chosen small enough.

Let \mathcal{C} be an arrangement of circles represented on the sphere. Each circle of \mathcal{C} spans a plane in 3-space, hence, we obtain an arrangement $\mathcal{E}(\mathcal{C})$ of planes in \mathbb{R}^3 . In fact, with a sphere S we get a bijection between (not necessarily connected) circle arrangements on S and arrangements of planes with the property that each plane of the arrangement intersects S .

Consider two circles C_1, C_2 of a circle arrangement \mathcal{C} on S and the corresponding planes E_1, E_2 of $\mathcal{E}(\mathcal{C})$. The intersection of E_1 and E_2 is either empty (i.e., E_1 and E_2 are parallel) or a line ℓ . The line ℓ intersects S if and only if C_1 and C_2 intersect, in fact, $\ell \cap S = C_1 \cap C_2$.

With three pairwise intersecting circles C_1, C_2, C_3 we obtain three planes E_1, E_2, E_3 intersecting in a vertex v of $\mathcal{E}(\mathcal{C})$. It is notable that v is in the interior of S if and only if the three circles form a Krupp in \mathcal{C} . We save this observation for further reference.

Fact 6.1. *Let \mathcal{C} be an arrangement of circles represented on the sphere. Three circles C_1, C_2, C_3 of \mathcal{C} form a Krupp if and only if the three corresponding planes E_1, E_2, E_3 intersect in a single point in the interior of S .*

For digons of \mathcal{C} , we also have necessary conditions in terms of $\mathcal{E}(\mathcal{C})$ and S .

Fact 6.2. *Let \mathcal{C} be an arrangement of circles represented on the sphere S . If a pair of intersecting circles C_1, C_2 in \mathcal{C} forms a digon of \mathcal{C} , then the line $E_1 \cap E_2$ has no intersection with any other plane E_3 corresponding to a circle $C_3 \notin \{C_1, C_2\}$ inside of S .*

6.1 Arrangements of Great-Pseudocircles

Central projections¹ – not to be confused with stereographic projections – map between arrangements of great-circles on a sphere S and arrangements of lines on a plane. Figure 6.1 gives an illustration. Changes of the plane to which we project preserve the isomorphism class of the projective arrangement of lines. In fact, arrangements of lines in the projective plane are in one-to-one correspondence to arrangements of great-circles.

¹recall the point-line duality from Part I of this thesis (cf. Figure 2.3)




Figure 6.1: Great-circle-line duality illustration. The points A, B, C, D from the sphere are projected to the points A', B', C', D' in the plane; the center of the sphere is also the center of projection.

In this section we generalize this concept to arrangements of pseudolines and show that there is a one-to-one correspondence to arrangements of great-pseudocircles. As already mentioned, this correspondence is not new (see e.g. [BLW⁺99]).

A *pseudoline* is a simple closed non-contractible curve in the projective plane. A (*projective*) *arrangement of pseudolines* is a collection of pseudolines such that any two intersect in exactly one point where they cross. We can also consider arrangements of pseudolines in the Euclidean plane by fixing a “line at infinity” in the projective plane – we call this a *projection*.

A Euclidean arrangement of n pseudolines can be represented by x -monotone pseudolines, a special representation of this kind is the wiring diagram ([Goo80], see also [FG18]). As illustrated in Figure 6.2, an x -monotone representation can be glued with a mirrored copy of itself to form an arrangement of n pseudocircles. The resulting arrangement is intersecting and has no NonKrupp subarrangement, hence, it is an arrangement of great-pseudocircles.




Figure 6.2: (a) Obtaining an arrangement \mathcal{A} of great-pseudocircles from an Euclidean arrangement \mathcal{L} of pseudolines and its mirrored copy. The gray boxes highlight the arrangement \mathcal{L} and its mirrored copy. (b) A great-circle representation of \mathcal{A} on the sphere.

For a pseudocircle C of an arrangement of n great-pseudocircles the cyclic order of crossings on C is *antipodal*, i.e., the infinite sequence corresponding to the cyclic order crossings of C with the other pseudocircles is periodic of order $n - 1$. If we consider projections of projective arrangements of n pseudolines, then this order does not depend on the choice of the projection.

In fact, projective arrangements of n pseudolines are in bijection with arrangements of n great-pseudocircles.

6.2 Incidence Theorems

The smallest non-stretchable arrangements of pseudolines are closely related to the incidence theorems of Pappos and Desargues. A construction already described by Levi [Lev26] is depicted in Figure 6.3(a). Pappos's Theorem states that, in a configuration of 8 lines as shown in the figure in black, the 3 white points are collinear, i.e., a line containing two of them also contains the third. Therefore, the arrangement including the red pseudoline has no corresponding arrangement of straight lines, i.e., it is not stretchable.

Miquel's Theorem asserts that, in a configuration of 5 circles as shown in Figure 6.3(b) in black, the 4 white points are cocircular, i.e., a circle containing three of them also contains the fourth. Therefore, the arrangement including the red pseudocircle cannot be circularized.




Figure 6.3: (a) A non-stretchable arrangement of pseudolines from Pappos's Theorem. (b) A non-circularizable arrangement of pseudocircles from Miquel's Theorem.

Next we state two incidence theorems that will be used in later proofs of non-circularizability. In the course of Chapter 7 we will meet further incidence theorems such as Lemma 7.11, Lemma 7.12, Theorem 7.14, Lemma 7.17, and again Miquel's Theorem (Theorem 7.18).

Lemma 6.1 (First Four-Circles Incidence Lemma). *Let \mathcal{C} be an arrangement of four circles C_1, C_2, C_3, C_4 such that none of them is contained in the interior of another one, and such that (C_1, C_2) , (C_2, C_3) , (C_3, C_4) , and (C_4, C_1) are touching. Then there is a circle C^* passing through these four touching points in the given cyclic order.*

We point the interested reader to the website “Cut-the-Knot.org” [Bog], where this lemma is stated (except for the cyclic order). The website also provides an interactive GeoGebra applet, which nicely illustrates the incidences.




Figure 6.4: An illustration for the proof of Lemma 6.1.

Proof. Apply a Möbius transformation Γ that maps the touching point of C_1 and C_2 to the point ∞ of the extended complex plane. This maps C_1 and C_2 to a pair L_1, L_2 of parallel lines. The discs of C_1 and C_2 are mapped to disjoint halfplanes. We may assume that L_1 and L_2 are horizontal and that L_1 is above L_2 . Circles C_3 and C_4 are mapped to touching circles C'_3 and C'_4 . Moreover, C'_3 is touching L_2 from above and C'_4 touches L_1 from below. Figure 6.4 shows a sketch of the situation.

Let ℓ' be the line, which is tangent to C'_3 and C'_4 at their touching point p . Consider the two segments from p to $C'_3 \cap L_2$ and from p to $C'_4 \cap L_1$. Elementary considerations show the following equalities of angles: $\alpha_3 = \alpha_4$, $\beta_3 = \gamma_3$, $\beta_4 = \gamma_4$, and $\gamma_3 = \gamma_4$ (cf. Figure 6.4). Hence, there is a line ℓ containing the images of the four touching points. Consequently, the circle $C^* = \Gamma^{-1}(\ell)$ contains the four touching points of \mathcal{C} , i.e., they are cocircular. \square

The following theorem (illustrated in Figure 6.5) is mentioned by Richter-Gebert [RG11, page 26] as a relative of Pappos's and Miquel's Theorem.

Theorem 6.2 ([RG11]). *Let C_1, C_2, C_3 be three circles in the plane such that each pair of them intersects in two points, and let ℓ_i be the line spanned by the two points of intersection of C_j and C_k , for $\{i, j, k\} = \{1, 2, 3\}$. Then ℓ_1, ℓ_2 , and ℓ_3 meet in a common point.*




Figure 6.5: An illustration of Theorem 6.2.

Proof. Use a stereographic projection ϕ to map the three circles to circles C'_1, C'_2, C'_3 on a sphere S . Consider the planes E'_1, E'_2, E'_3 spanned by C'_1, C'_2, C'_3 . Let ℓ'_i be the line $E'_j \cap E'_k$, for $\{i, j, k\} = \{1, 2, 3\}$. Since the arrangement is simple and intersecting, the lines $\ell'_1, \ell'_2, \ell'_3$ are distinct and the intersection $E'_1 \cap E'_2 \cap E'_3$ is a single projective point p , which is contained in each of $\ell'_1, \ell'_2, \ell'_3$. The inverse of ϕ can be interpreted as a central projection from 3-space to the plane. In this interpretation of ϕ^{-1} , the lines $\ell'_1, \ell'_2, \ell'_3$ are mapped to ℓ_1, ℓ_2, ℓ_3 and p is mapped to a projective point, i.e., either p is a point or the lines are parallel (i.e. p lies at infinity). \square

6.3 Flips and Deformations

Let \mathcal{C} be an arrangement of circles. Imagine that the circles of \mathcal{C} start moving independently, i.e., the position of their centers and their radii depend on a time parameter t in a continuous way. This yields a family $\mathcal{C}(t)$ of arrangements with $\mathcal{C}(0) = \mathcal{C}$. Let us assume that the set T of all t for which $\mathcal{C}(t)$ is not simple or contains touching circles is discrete and for each $t \in T$ the arrangement $\mathcal{C}(t)$ contains either a single point where 3 circles intersect or a single touching. If $t_1 < t_2$ are consecutive in T , then all arrangements $\mathcal{C}(t)$ with $t \in (t_1, t_2)$ are isomorphic. Selecting one representative from each such class, we get a list $\mathcal{C}_0, \mathcal{C}_1, \dots$ of simple arrangements

such that two consecutive (non-isomorphic) arrangements $\mathcal{C}_i, \mathcal{C}_{i+1}$ are either related by a triangle flip or by a digon flip, see Figure 6.6.




Figure 6.6: An illustration of the flip operations.

We will make use of controlled changes in circle arrangements, in particular, we grow or shrink specified circles of an arrangement to produce touchings or points where 3 circles intersect. The following lemma will be of use frequently.

Lemma 6.3 (Digon Collapse Lemma). *Let \mathcal{C} be an arrangement of circles in the plane and let C be one of the circles of \mathcal{C} , which intersects at least two other circles from \mathcal{C} and does not fully contain any other circle from \mathcal{C} in its interior. If C has no incident triangle in its interior, then we can shrink C into its interior such that the combinatorics of the arrangement remain the same except that two digons collapse to touchings. Moreover, the two corresponding circles touch C from the outside.*




Figure 6.7: An illustration of the Digon Collapse Lemma.

Proof. As illustrated on the left hand side of Figure 6.7, we shrink the radius of C until the first event occurs. Since C does not fully contain any other circle from \mathcal{C} in its interior, no new digon can be created. Moreover, since C has no incident triangles in its interior, an interior digon collapses. We obtain a point where C touches another circle that lies outside of C . (Note that several digons might collapse at the same time.)

If C has only one touching point p , we shrink the radius and simultaneously move the center towards p (cf. the right hand side of Figure 6.7) such that p stays a touching until a second digon becomes a touching. Again the touching point is with a circle that lies outside of C . \square

In the following we will sometimes use the dual version of the lemma, whose statement is obtained from the Digon Collapse Lemma by changing interior to exterior and outside to inside. The validity of the dual lemma is seen by applying a Möbius transformation which exchanges interior and exterior of C .

Triangle flips and digon flips are also central to the work of Snoeyink and Hershberger [SH91]. They have shown that an arrangement \mathcal{C} of pseudocircles can be swept with a sweepfront γ starting at any pseudocircle $C \in \mathcal{C}$, i.e., $\gamma_0 = C$. The sweep consists of two stages, one for sweeping the interior of C , the other for sweeping the exterior. At any fixed time t the sweepfront γ_t is a closed curve such that $\mathcal{C} \cup \{\gamma_t\}$ is an arrangement of pseudocircles. Moreover, this

arrangement is simple except for a discrete set T of times where sweep events happen. The sweep events are triangle flips or digon flips involving γ_t .

Chapter 7

On Circularizability

In Chapter 7.1 we present one of our main results of this thesis:

Theorem 7.1 (Great-Circle Theorem). *An arrangement of great-pseudocircles is circularizable (i.e., has a circle representation) if and only if it has a great-circle representation.*

The Great-Circle Theorem allows to transfer knowledge regarding arrangements of pseudolines to arrangements of pseudocircles. Subsequent to the theorem, we present several direct consequences such as the $\exists\mathbb{R}$ -completeness of circularizability.

In Chapters 7.2 and 7.3 respectively, we present the full classification of circularizable and non-circularizable arrangements among all connected arrangements of 5 pseudocircles and all digon-free intersecting arrangements of 6 pseudocircles. With the aid of computers we generated the complete lists of connected arrangements of $n \leq 6$ pseudocircles and of intersecting arrangements of $n \leq 7$ pseudocircles. For the class of arrangements of n great-pseudocircles, the numbers were already determined for $n \leq 11$ [Knu92, AAK02, Kra03, AK06]. The respective numbers are shown in Table 7.1 (cf. sequences A288568, A296406, and A006248 on the OEIS [Slo]). Given the complete lists of arrangements, we used automatized heuristics to find circle representations. Examples where the heuristics failed had to be examined by hand.

n	3	4	5	6	7	8
connected	3	21	984	609 423	?	
+digon-free	1	3	30	4 509	?	
connected+cylindrical	3	20	900	530 530	?	
+digon-free	1	3	30	4 477	?	
intersecting	2	8	278	145 058	447 905 202	?
+digon-free	1	2	14	2 131	3 012 972	?
intersecting+cylindrical	2	8	278	144 395	436 634 633	?
+digon-free	1	2	14	2 131	3 012 906	?
great-pseudocircles	1	1	1	4	11	135

n	9	10	11
great-pseudocircles	4 382	312 356	41 848 591

Table 7.1: Number of combinatorially different arrangements of n pseudocircles.

Computational issues and algorithmic ideas are deferred until Chapter 7.5. There we also sketch the heuristics that we have used to produce circle representations for most of the arrangements. The encoded lists of arrangements of up to $n = 6$ pseudocircles and circle representations

are available on our webpage [FS]. Chapter 7.5 also contains asymptotic results on the number of arrangements of n pseudocircles as well as results on their flip-graph.

The list of circle representations at [FS] together with the non-circularizability proofs given in Chapter 7.2 yields the following theorem.

Theorem 7.2. *The four isomorphism classes of arrangements \mathcal{N}_5^1 , \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4 (shown in Figure 7.1) are the only non-circularizable ones among the 984 isomorphism classes of connected arrangements of $n = 5$ pseudocircles.*

Corollary 7.3. *The isomorphism class of arrangement \mathcal{N}_5^1 is the unique non-circularizable one among the 278 isomorphism classes of intersecting arrangements of $n = 5$ pseudocircles.*

We remark that the arrangements \mathcal{N}_5^1 , \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4 have symmetry groups of order 4, 8, 2, and 4, respectively. Also, note that none of the four examples is digon-free. Non-circularizability of \mathcal{N}_5^2 was previously shown by Linhart and Ortner [LO05]. We give an alternative proof which also shows the non-circularizability of \mathcal{N}_5^3 . Jonathan Wild and Christopher Jones, contributed sequences A250001 and A288567 to the On-Line Encyclopedia of Integer Sequences (OEIS) [Slo]. These sequences count certain classes of arrangements of circles and pseudocircles. Wild and Jones also looked at circularizability and independently found Theorem 7.2 (personal communication).




Figure 7.1: The four non-circularizable arrangements on $n = 5$ pseudocircles: (a) \mathcal{N}_5^1 , (b) \mathcal{N}_5^2 , (c) \mathcal{N}_5^3 , and (d) \mathcal{N}_5^4 .

Concerning arrangements of 6 pseudocircles, we were able to fully classify digon-free intersecting arrangements.

Theorem 7.4. *The three isomorphism classes of arrangements \mathcal{N}_6^Δ , \mathcal{N}_6^2 , and \mathcal{N}_6^3 (shown in Figure 7.2) are the only non-circularizable ones among the 2131 isomorphism classes of digon-free intersecting arrangements of $n = 6$ pseudocircles.*

In Chapter 7.3, we give non-circularizability proofs for \mathcal{N}_6^Δ , \mathcal{N}_6^2 , and \mathcal{N}_6^3 . In fact, for the non-circularizability of \mathcal{N}_6^Δ and \mathcal{N}_6^2 , respectively, we have two proofs of different flavors: One proof uses continuous deformations similar to the proof of the Great-Circle Theorem (Theorem 7.1) and the other proof is based on an incidence theorem. The incidence theorem used for \mathcal{N}_6^Δ may be of independent interest:

Theorem 7.5. *Let a, b, c, d, A, B, C, D be 8 points from \mathbb{R}^3 such that no five of these points lie in a common plane and each of the following 5 subsets of 4 points is coplanar:*

$$\{a, b, A, B\}, \{a, c, A, C\}, \{a, d, A, D\}, \{b, c, B, C\}, \text{ and } \{b, d, B, D\}.$$

Then $\{c, d, C, D\}$ is also coplanar.

The proof we give is based on determinant cancellation, a technique that we learned from Richter-Gebert, cf. [RG11]. It turned out that this incidence theorem is a slight generalization of the *Bundle theorem* (see e.g. [Wika]), where the points are typically assumed to lie on the sphere.




Figure 7.2: The three non-circularizable digon-free intersecting arrangements for $n = 6$: (a) \mathcal{N}_6^Δ , (b) \mathcal{N}_6^2 , and (c) \mathcal{N}_6^3 . Inner triangles are colored gray. Note that in (b) and (c) the outer face is a triangle.

An instance of Theorem 7.5 is obtained by assigning the eight letters appropriately to the corners of a cube. Each of the six involved sets then corresponds to the four corners covered by two opposite edges of the cube. The appropriate assignment of the letters can be derived from Figure 7.9.

We remark that the arrangements \mathcal{N}_6^Δ , \mathcal{N}_6^2 , and \mathcal{N}_6^3 have symmetry groups of order 24, 3, and 6, respectively. Particularly interesting is the arrangement \mathcal{N}_6^Δ (Figure 7.2(a), see also Figure 7.9). This is the unique intersecting digon-free arrangement of 6 pseudocircles which attains the minimum 8 for the number of triangles (see [FS19b]).

Even though we could not complete the classification for intersecting arrangements of 6 pseudocircles, we provide some further nice incidence theorems and non-circularizability proofs in Chapter 7.4. One of them is the example of Edelsbrunner and Ramos [ER97]; see Figure 7.14(a) and Figure 5.5(left).

It may be worth mentioning that, by enumerating and realizing all arrangements of $n \leq 4$ pseudocircles, we have an alternative proof of the Kang and Müller result, that all arrangements of $n \leq 4$ pseudocircles are circularizable [KM14].

7.1 The Great-Circle Theorem and its Applications

Let \mathcal{A} be an arrangement of great-pseudocircles and let \mathcal{L} be the corresponding projective arrangement of pseudolines (cf. Chapter 6.1). Central projections show that, if \mathcal{L} is realizable with straight lines, then \mathcal{A} is realizable with great-circles, and conversely. In fact, due to Theorem 7.1, it is sufficient that \mathcal{A} is circularizable to conclude that \mathcal{A} is realizable with great-circles and \mathcal{L} is realizable with straight lines.

Theorem 7.1 (Great-Circle Theorem). *An arrangement of great-pseudocircles is circularizable (i.e., has a circle representation) if and only if it has a great-circle representation.*

Proof of Theorem 7.1. Consider an arrangement of circles \mathcal{C} on the unit sphere \mathbb{S} that realizes an arrangement of great-pseudocircles as illustrated in Figure 7.3(left and center). Let $\mathcal{E}(\mathcal{C})$ be the arrangement of planes spanned by the circles of \mathcal{C} . Since \mathcal{C} realizes an arrangement of great-pseudocircles, every triple of circles forms a Krupp, hence, the point of intersection of any three planes of $\mathcal{E}(\mathcal{C})$ is in the interior of \mathbb{S} .

Imagine the planes of $\mathcal{E}(\mathcal{C})$ moving towards the origin. To be precise, for time $t \geq 1$ let $\mathcal{E}_t := \{1/t \cdot E : E \in \mathcal{E}(\mathcal{C})\}$. Since all intersection points of the initial arrangement $\mathcal{E}_1 = \mathcal{E}(\mathcal{C})$ are in the interior of the unit sphere \mathbb{S} , the circle arrangement obtained by intersecting the moving planes \mathcal{E}_t with \mathbb{S} remains the same (isomorphic). Moreover, as illustrated in Figure 7.3(right),

every circle in this arrangement converges to a great-circle as $t \rightarrow +\infty$, and the statement follows. \square




Figure 7.3: Illustration for the proof of Theorem 7.1.

The Great-Circle Theorem has several interesting consequences. The following corollary allows us to transfer results from the world of pseudolines into the world of (great-)pseudocircles.

Corollary 7.6. *An arrangement of pseudolines is stretchable if and only if the corresponding arrangement of great-pseudocircles is circularizable.*

Since deciding stretchability of arrangements of pseudolines is known to be ETR-complete (see e.g. [Mnë88, Mat14, ŠS17]), the hardness of stretchability directly carries over to hardness of circularizability. To show containment in ETR, the circularizability problem has to be modeled with polynomial inequalities. This can be done by taking the centers and radii of the circles as variables and using polynomial inequalities to prescribe the order of the intersections along the circles. A detailed description of such polynomial inequalities is deferred to Chapter 7.5.3.

Corollary 7.7. *Deciding circularizability is ETR-complete, even when the input is restricted to arrangements of great-pseudocircles.*

It is known that all (not necessarily simple) arrangements of $n \leq 8$ pseudolines are stretchable and that the simple non-Pappos arrangement is the unique non-stretchable simple projective arrangement of 9 pseudolines, see e.g. [FG18]. This again carries over to arrangements of great-pseudocircles.

Corollary 7.8. *All arrangements of up to 8 great-pseudocircles are circularizable and the arrangement corresponding to the simple non-Pappos arrangement of pseudolines is the unique non-circularizable arrangement of 9 great-pseudocircles.*

Note that the statement of the corollary also holds for the non-simple case. A non-simple arrangement of great-pseudocircles is an arrangement where three pseudocircles either form a Krupp or the intersection of the three pseudocircles consists of two points. Grünbaum [Grü80] denoted arrangements of great-pseudocircles as “symmetric”.

Bokowski and Sturmfels [BS89] have shown that infinite families of minimal non-stretchable arrangements of pseudolines exist, i.e., non-stretchable arrangements where every proper subarrangement is stretchable. Again, this carries over to arrangements of pseudocircles.

Corollary 7.9. *There exist infinite families of minimal non-circularizable arrangements of (great-)pseudocircles.*

Mnëv’s Universality Theorem [Mnë88], see also [RG95], has strong implications for pseudoline arrangements and stretchability. Besides the hardness of stretchability, it also shows the existence of arrangements of pseudolines with a disconnected realization space, that is, there are isomorphic arrangements of lines such that there is no continuous transformation which transforms one arrangement into the other within the isomorphism class. Suvorov [Suv88] gave an explicit example of two such arrangements on $n = 13$ lines.

Corollary 7.10. *There exist circularizable arrangements of (great-)pseudocircles with a disconnected realization space.*

7.2 Arrangements of 5 Pseudocircles

In this section we prove Theorem 7.2.

Theorem 7.2. *The four isomorphism classes of arrangements \mathcal{N}_5^1 , \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4 (shown in Figure 7.1) are the only non-circularizable ones among the 984 isomorphism classes of connected arrangements of $n = 5$ pseudocircles.*

On the webpage [FS] we have the data for circle realizations of 980 out of the 984 connected arrangements of 5 pseudocircles. The remaining four arrangements in this class are the four arrangements of Theorem 7.2. Since all arrangements with $n \leq 4$ pseudocircles have circle representations, there are no disconnected non-circularizable examples with $n \leq 5$. Hence, the four arrangements \mathcal{N}_5^1 , \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4 are the only non-circularizable arrangements with $n \leq 5$. Since \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4 are not intersecting, \mathcal{N}_5^1 is the unique non-circularizable intersecting arrangement of 5 pseudocircles, this is Corollary 7.3.

7.2.1 Non-circularizability of \mathcal{N}_5^1

The arrangement \mathcal{N}_5^1 is depicted in Figures 7.1(a) and 7.5. Since Figure 7.1(a) is meant to illustrate the symmetry of \mathcal{N}_5^1 while Figure 7.5 illustrates our non-circularizability proof, the two drawings of \mathcal{N}_5^1 differ in the Euclidean plane. However, we have marked one of the cells of \mathcal{N}_5^1 by black cross in both figures to highlight the isomorphism.

For the non-circularizability proof of \mathcal{N}_5^1 we will make use of the following incidence lemma.

Lemma 7.11 (Second Four-Circles Incidence Lemma). *Let \mathcal{C} be an arrangement of four circles C_1, C_2, C_3, C_4 such that every pair of them is touching or forms a digon in \mathcal{C} , and every circle is involved in at least two touchings. Then there is a circle C^* passing through the digon or touching point of each of the following pairs of circles (C_1, C_2) , (C_2, C_3) , (C_3, C_4) , and (C_4, C_1) in this cyclic order.*




Figure 7.4: Illustration for the proof of Lemma 7.11.

Proof. We first deal with the case where C_1 and C_3 form a digon. The assumptions imply that there is at most one further digon which might then be formed by C_2 and C_4 . In particular, the four pairs mentioned in the statement of the lemma form touchings and, as illustrated in the first row of Figure 7.4, we will find a circle C^* that is incident to those four touching points. In the following let p_{ij} denote the touching point of C_i and C_j .

Think of the circles as being in the extended complex plane. Apply a Möbius transformation Γ that maps one of the points of intersection of C_1 and C_3 to the point ∞ . This maps C_1 and C_3 to a pair of crossing lines. The images of C_2 and C_4 are circles which touch the two lines corresponding to C_1 and C_3 . The first row of Figure 7.4 gives an illustration. Since the centers of C_2 and C_4 lie on the bisector ℓ of the lines $\Gamma(C_1)$ and $\Gamma(C_3)$, the touchings of C_2 and C_4 are symmetric with respect to ℓ . Therefore, there is a circle C with center on ℓ that contains the images of the four points p_{12} , p_{23} , p_{34} , and p_{41} . The circle $C^* = \Gamma^{-1}(C)$ contains the four points, i.e., they are cocircular.

Now we consider the case where C_1 and C_3 touch. Again we apply a Möbius transformation Γ that sends p_{13} to ∞ . This maps C_1 and C_3 to parallel lines, each touched by one of C_2 and C_4 . The second row of Figure 7.4 shows that there is a circle C such that $C^* = \Gamma^{-1}(C)$ has the claimed property. \square




Figure 7.5: An illustration of the non-circularizability proof of \mathcal{N}_5^1 . The auxiliary circle C^* is drawn dashed.

Proof (non-circularizability of \mathcal{N}_5^1). Suppose for a contradiction that there is an isomorphic arrangement \mathcal{C} of circles. We label the circles as illustrated in Figure 7.5 and apply the Digon Collapse Lemma (Lemma 6.3) to shrink C_2 , C_3 , and C_4 into their respective interiors. We also use the dual of the Digon Collapse Lemma for C_1 . In the resulting subarrangement \mathcal{C}' formed by these four transformed circles C'_1, C'_2, C'_3, C'_4 , each of the four circles is involved in at least two touchings. By applying Lemma 7.11 to \mathcal{C}' we obtain a circle C^* which passes through the four points p_{12} , p_{23} , p_{34} , and p_{41} (in this order) which respectively are touching points or points from the digons of (C'_1, C'_2) , (C'_2, C'_3) , (C'_3, C'_4) , and (C'_4, C'_1) .

Moreover, since the intersection of C'_i and C'_j in \mathcal{C}' is contained in the intersection of C_i and C_j in \mathcal{C} , each of the four points p_{12} , p_{23} , p_{34} , and p_{41} lies in the original digon of \mathcal{C} . It follows that the circle C_5 has p_{12} and p_{34} in its interior but p_{23} and p_{41} in its exterior (cf. Figure 7.5). By applying Lemma 7.11 to \mathcal{C}' we obtain a circle C^* which passes through the points p_{12} , p_{23} , p_{34} , and p_{41} (in this order). Now the two circles C_5 and C^* intersect in four points. This is impossible, and hence \mathcal{N}_5^1 is not circularizable. \square

7.2.2 Non-circularizability of the connected arrangements \mathcal{N}_5^2 , \mathcal{N}_5^3 , and \mathcal{N}_5^4

The non-circularizability of \mathcal{N}_5^2 has been shown by Linhart and Ortner [LO05]. We give an alternative proof which also shows the non-circularizability of \mathcal{N}_5^3 . The two arrangements \mathcal{N}_5^2 and \mathcal{N}_5^3 are depicted in Figures 7.1(b) and 7.1(c), and also in Figures 7.6(a) and 7.6(b).




Figure 7.6: An illustration of the non-circularizability proofs of (a) \mathcal{N}_5^2 and (b) \mathcal{N}_5^3 . The auxiliary circle C^* is drawn dashed.

Proof (non-circularizability of \mathcal{N}_5^2 and \mathcal{N}_5^3). Suppose for a contradiction that there is an isomorphic arrangement \mathcal{C} of circles. We label the circles as illustrated in Figure 7.6. Since the respective interiors of C_1 and C_3 are disjoint, we can apply the Digon Collapse Lemma (Lemma 6.3) to C_1 and C_3 . This yields an arrangement \mathcal{C}' with four touching points $p_{12}, p_{23}, p_{34}, p_{41}$, where p_{ij} is the touching point of C'_i and C'_j .

From Lemma 6.1 it follows that there is a circle C^* which passes through the points p_{12}, p_{23}, p_{34} , and p_{41} in this cyclic order. Since the point p_{ij} lies inside the digon formed by C_i and C_j in the arrangement \mathcal{C} , it follows that the circle C_5 has p_{12}, p_{34} in its interior and p_{23}, p_{41} in its exterior. Therefore, the two circles C_5 and C^* intersect in four points. This is impossible and, therefore, \mathcal{N}_5^2 and \mathcal{N}_5^3 are not circularizable. \square

It remains to prove that \mathcal{N}_5^4 (shown in Figures 7.1(d) and 7.8) is not circularizable. In the proof we make use of the following incidence lemma.

Lemma 7.12 (Third Four-Circles Incidence Lemma). *Let \mathcal{C} be an arrangement of four circles C_1, C_2, C_3, C_4 such that (C_1, C_2) , (C_2, C_3) , (C_4, C_1) , and (C_3, C_4) are touching, moreover, C_4 is in the interior of C_1 and the exterior of C_3 , and C_2 is in the interior of C_3 and the exterior of C_1 , see Figure 7.7. Then there is a circle C^* passing through the four touching points in the given cyclic order.*

Proof. Since C_1 touches C_2 and C_4 which are respectively inside and outside C_3 , the two circles C_1 and C_3 intersect. Apply a Möbius transformation Γ that maps a crossing point of C_1 and C_3 to the point ∞ of the extended complex plane. This maps C_1 and C_3 to a pair L_1, L_3 of lines. The images C'_2, C'_4 of C_2 and C_4 are separated by the lines L_1, L_3 and each of them touches both lines. Figure 7.7 illustrates the situation. The figure also shows that a circle C' through the four touching points exists. The circle $C^* = \Gamma^{-1}(C')$ has the claimed properties. \square

Proof (non-circularizability of \mathcal{N}_5^4). Suppose for a contradiction that there is an isomorphic arrangement \mathcal{C} of circles. We label the circles as illustrated in Figure 7.8. We shrink the circles C_2 and C_4 such that each of the pairs (C_1, C_2) , (C_2, C_3) , (C_3, C_4) , and (C_4, C_1) touch. (Note that




Figure 7.7: An illustration for the proof of Lemma 7.12.




Figure 7.8: Illustration of the non-circularizability proof of the arrangement \mathcal{N}_5^4 . The circle C^* is drawn dashed.

each of these pairs forms a digon in \mathcal{C} .) With these touchings the four circles C_1, C_2, C_3, C_4 form the configuration of Lemma 7.12. Hence there is a circle C^* containing the four touching points in the given cyclic order. Now the two circles C^* and C_5 intersect in four points. This is impossible and, therefore, \mathcal{N}_5^4 is not circularizable. \square

7.3 Intersecting Digon-free Arrangements of 6 Pseudocircles

In this section we prove Theorem 7.4.

Theorem 7.4. *The three isomorphism classes of arrangements \mathcal{N}_6^Δ , \mathcal{N}_6^2 , and \mathcal{N}_6^3 (shown in Figure 7.2) are the only non-circularizable ones among the 2131 isomorphism classes of digon-free intersecting arrangements of $n = 6$ pseudocircles.*

We remark that all three arrangements do not have the intersecting arrangement \mathcal{N}_5^1 as a subarrangement – otherwise non-circularizability would follow directly. In fact, \mathcal{N}_5^1 has no extension to an intersecting digon-free arrangement of six pseudocircles.

On the webpage [FS] we have the data of circle realizations of all 2131 intersecting digon-free arrangements of 6 pseudocircles except for the three arrangements mentioned in Theorem 7.4. In the following, we present two non-circularizability proofs for \mathcal{N}_6^Δ and \mathcal{N}_6^2 , respectively, and a non-circularizability proof for \mathcal{N}_6^3 .

7.3.1 Non-circularizability of \mathcal{N}_6^Δ

The arrangement \mathcal{N}_6^Δ (shown in Figure 7.2(a)) is an intersecting digon-free arrangement. Our interest in \mathcal{N}_6^Δ was originally motivated by our study [FS19b] of arrangements of pseudocircles with few triangles. From a computer search we know that \mathcal{N}_6^Δ occurs as a subarrangement

of every digon-free arrangement for $n = 7, 8, 9$ with $p_3 < 2n - 4$ triangles. Since \mathcal{N}_6^Δ is not circularizable, neither are these arrangements. It thus seems plausible that every arrangement of n circles has at least $p_3 \geq 2n - 4$ triangles – but we will come back to this later in Chapter 8 (cf. Conjecture 8.1).

Our first proof is an immediate consequence of the following theorem, whose proof resembles the proof of the Great-Circle Theorem (Theorem 7.1).

Theorem 7.13. *Let \mathcal{A} be a connected digon-free arrangement of pseudocircles. If every triple of pseudocircles which forms a triangle is NonKrupp, then \mathcal{A} is not circularizable.*

Proof. Assume for a contradiction that there exists an isomorphic arrangement of circles \mathcal{C} on the unit sphere \mathbb{S} . Let $\mathcal{E}(\mathcal{C})$ be the arrangements of planes spanned by the circles of \mathcal{C} .

Imagine the planes of $\mathcal{E}(\mathcal{C})$ moving away from the origin. To be precise, for time $t \geq 1$ let $\mathcal{E}_t := \{t \cdot E : E \in \mathcal{E}(\mathcal{C})\}$. Consider the arrangement induced by intersecting the moving planes \mathcal{E}_t with the unit sphere \mathbb{S} . Since \mathcal{C} has NonKrupp triangles, it is not a great-circle arrangement and some planes of $\mathcal{E}(\mathcal{C})$ do not contain the origin. All planes from $\mathcal{E}(\mathcal{C})$, which do not contain the origin, will eventually lose the intersection with \mathbb{S} , hence some event has to happen.

When the isomorphism class of the intersection of \mathcal{E}_t with \mathbb{S} changes, we see a triangle flip, or a digon flip, or some isolated circle disappears. Since initially there is no digon and no isolated circle, the first event is a triangle flip. By assumption, triangles of \mathcal{C} correspond to NonKrupp subarrangements, hence, the intersection point of their planes is outside of \mathbb{S} (Fact 6.1). This shows that a triangle flip event is also impossible. This contradiction implies that \mathcal{A} is non-circularizable. \square

Proof (first proof of non-circularizability of \mathcal{N}_6^Δ). The arrangement \mathcal{N}_6^Δ is intersecting, digon-free, and each of the eight triangles of \mathcal{N}_6^Δ is formed by three circles which are a NonKrupp configuration. Hence, Theorem 7.13 implies that \mathcal{N}_6^Δ is not circularizable. \square

All arrangements known to us whose non-circularizability can be shown with Theorem 7.13 contain \mathcal{N}_6^Δ as a subarrangement – which already shows non-circularizability. Based on this data we venture the following conjecture:

Conjecture 7.1. *Every connected digon-free arrangement \mathcal{A} of pseudocircles with the property, that every triple of pseudocircles which forms a triangle in \mathcal{A} is NonKrupp, contains \mathcal{N}_6^Δ as a subarrangement.*

Our second proof of non-circularizability of \mathcal{N}_6^Δ is based on an incidence theorem for circles (Theorem 7.14) which is a consequence of an incidence theorem for points and planes in 3-space (Theorem 7.5). Before going into details, let us describe the geometry of the arrangement \mathcal{N}_6^Δ : Consider the non-simple arrangement \mathcal{A}^\bullet obtained from \mathcal{N}_6^Δ by contracting each of the eight triangles into a single point of triple intersection. The arrangement \mathcal{A}^\bullet is circularizable. A realization is obtained by taking a cube inscribed in the sphere S such that each of the eight corners touches the sphere S . The arrangement \mathcal{A}^\bullet is the intersection of S with the six planes which are spanned by pairs of diagonally opposite edges of the cube.

Theorem 7.5. *Let a, b, c, d, A, B, C, D be 8 points from \mathbb{R}^3 such that no five of these points lie in a common plane and each of the following 5 subsets of 4 points is coplanar:*

$$\{a, b, A, B\}, \{a, c, A, C\}, \{a, d, A, D\}, \{b, c, B, C\}, \text{ and } \{b, d, B, D\}.$$

Then $\{c, d, C, D\}$ is also coplanar.

Proof. By assumption, the five points a, b, c, d, D do not lie in a common plane. If we assume – towards a contradiction – that the three points a, b, c lie on a common line, then c lies in the plane spanned by $\{a, b, A\}$. By assumption, the point B also lies in that plane, which then

contains the five points $\{a, b, A, B, c\}$ – a contradiction. Hence, the three points a, b, c are not collinear.

Now an easy case distinction shows that at most one of the 4-element subsets of the five points a, b, c, d, D is coplanar. Since the roles of d and D can be exchanged, we assume without loss of generality that a, b, c, d are affinely independent.

We now embed \mathbb{R}^3 as the hyperplane $\sum x_i = 1$ into \mathbb{R}^4 such that the four points a, b, c, d become the elements of the standard basis, namely, $a = \mathbf{e}_1$, $b = \mathbf{e}_2$, $c = \mathbf{e}_3$, and $d = \mathbf{e}_4$.

Now coplanarity of 4 points can be tested by evaluating the determinant. Coplanarity of $\{a, b, A, B\}$ yields $\det[abAB] = \det(\mathbf{e}_1, \mathbf{e}_2, A, B) = 0$. On the basis of the 5 collinear sets, we get the following determinants and equations:

$$\begin{array}{ccccc} \det[abAB] = 0 & \det[caCA] = 0 & \det[adAD] = 0 & \det[cbCB] = 0 & \det[bdBD] = 0 \\ A_3B_4 = A_4B_3 & A_4C_2 = A_2C_4 & A_2D_3 = A_3D_2 & B_1C_4 = B_4C_1 & B_3D_1 = B_1D_3 \end{array} .$$

Take the product of the left sides of the five equations and the product of the right sides. These products are the same. We can cancel as much as possible from the resulting equations and obtain $C_2D_1 = C_1D_2$. This cancellation works because all relevant factors are non-zero, which is shown below. We then obtain that $\det(\mathbf{e}_3, \mathbf{e}_4, C, D) = 0$, i.e., the coplanarity of $\{c, d, C, D\}$.

To show that $A_3 \neq 0$, suppose that $A_3 = 0$. This implies that $\{a, b, d, A\}$ is coplanar. Suppose that a, b, A are collinear. Since $\{a, b, d, A\}$ and $\{a, d, A, D\}$ are coplanar, the five points $\{a, b, d, A, D\}$ are coplanar – which is a contradiction. Hence a, b, A are not collinear. Now, since $\{a, b, A, d\}$ and $\{a, b, A, B\}$ are coplanar, the five points $\{a, b, A, B, d\}$ are coplanar – a contradiction to the assumption $A_3 = 0$. The proof that the other relevant factors are non-zero is analogous. \square

The theorem implies the following incidence theorem for circles.

Theorem 7.14. *Let C_1, C_2, C_3, C_4 be four circles and let a, b, c, d, w, x, y, z be eight distinct points in \mathbb{R}^2 such that $C_1 \cap C_2 = \{a, w\}$, $C_3 \cap C_4 = \{b, x\}$, $C_1 \cap C_3 = \{c, y\}$, and $C_2 \cap C_4 = \{d, z\}$. If there is a circle C containing a, b, w, x , then there is a circle C' containing c, d, y, z . Moreover, if one of the triples of C_1, C_2, C_3, C_4 forms a Krupp, then c, d, y, z represents the circular order on C' .*

Proof. Consider the arrangement of circles on the sphere. The idea is to apply Theorem 7.5. The coplanarity of the 5 sets follows because the respective 4 points belong to C, C_1, C_2, C_3, C_4 in this order. The 5 points a, b, c, d, z do not lie in a common plane because otherwise we would have $C_2 = C_4$ (a contradiction to $C_2 \cap C_4 = \{d, z\}$). Analogous arguments show that also none of the other 5-element subset of $\{a, b, c, d, x, y, z, w\}$ is coplanar. This shows that Theorem 7.5 can be applied. Regarding the circular order on C' , suppose that C_1, C_2, C_3 is a Krupp. This implies that C_2 separates c and y . Since $C' \cap C_2 = \{d, z\}$ the points of $\{c, y\}$ and $\{d, z\}$ alternate on C' , this implies the claim. \square

It is worth mentioning, that the second part of the theorem can be strengthened: If one of the triples of C_1, C_2, C_3, C_4 forms a Krupp, then the arrangement together with C and C' is isomorphic to the simplicial arrangement \mathcal{A}^\bullet obtained from \mathcal{N}_6^Δ by contracting the eight triangles into triple intersections. A *simplicial arrangement* is a non-simple arrangement where all cells are triangles. The arrangement \mathcal{A}^\bullet can be extended to larger simplicial arrangements by adding any subset of the three circles C_1^*, C_2^*, C_3^* which are defined as follows: C_1^* is the circle through the four points $(C_1 \cap C_4) \cup (C_2 \cap C_3)$; C_2^* is the circle through the four points $(C_1 \cap C_4) \cup (C \cap C')$; C_3^* is the circle through the four points $(C_2 \cap C_3) \cup (C \cap C')$. In each case the cocircularity of the four points defining C_i^* is a consequence of the theorem.

The following lemma is similar to the Digon Collapse Lemma (Lemma 6.3). By changing interior to exterior and outside to inside (by applying a Möbius transformation), we obtain a dual version also for this lemma.

Lemma 7.15 (Triangle Collapse Lemma). *Let \mathcal{C} be an arrangement of circles in the plane and let C be one of the circles of \mathcal{C} , which intersects at least three other circles from \mathcal{C} and does not fully contain any other circle from \mathcal{C} in its interior. If C has no incident digon in its interior, then we can continuously transform C such that the combinatorics of the arrangement remain the same except that two triangles collapse to points of triple intersection. Moreover, it is possible to prevent a fixed triangle T incident to C from being the first one to collapse.*

Proof. The proof is very much like the proof of Lemma 6.3. Shrink the radius of C until the first flip occurs, this must be a triangle flip, i.e., a triangle is reduced to a point of triple intersection. If C has a point p of triple intersection, shrink C towards p , i.e., shrink the radius and simultaneously move the center towards p such that p stays incident to C . With the next flip a second triangle collapses.

For the extension let $q \in T \cap C$ be a point. Start the shrinking process by shrinking C towards q . This prevents T from collapsing. \square




Figure 7.9: The arrangement \mathcal{N}_6^Δ with a labeling of its eight triangles.

Proof (second proof of non-circularizability of \mathcal{N}_6^Δ). Suppose for a contradiction that \mathcal{N}_6^Δ has a realization \mathcal{C} . Each circle of \mathcal{C} has exactly two incident triangles in the inside and exactly two on the outside. Apply Lemma 7.15 to C_5 and to C_3 (we refer to the circles with the colors and labels used in Figure 7.9). This collapses the triangles labeled a, b, x, w , i.e., all the triangles incident to C_5 . Now the green circle C_1 , the magenta circle C_2 , the black circle C_3 , the blue circle C_4 , and the red circle C_5 are precisely the configuration of Theorem 7.14 with C_5 in the role of C . The theorem implies that there is a circle C' containing the green-black crossing at c , the blue-magenta crossing at d , the green-black crossing at y , and the blue-magenta crossing at z in this order. Each consecutive pair of these crossings is on different sides of the yellow circle C_6 , hence, there are at least four crossings between C' and C_6 . This is impossible for circles, whence, there is no circle arrangement \mathcal{C} realizing \mathcal{N}_6^Δ . \square

7.3.2 Non-circularizability of \mathcal{N}_6^2

The arrangement \mathcal{N}_6^2 is shown in Figure 7.2(b) and Figure 7.11(a). We give two proofs for the non-circularizability of \mathcal{N}_6^2 . The first one is an immediate consequence of the following theorem, which – in the same flavor as Theorem 7.13 – can be obtained similarly as the proof of the Great-Circle Theorem (Theorem 7.1).

Theorem 7.16. *Let \mathcal{A} be an intersecting arrangement of pseudocircles which is not an arrangement of great-pseudocircles. If every triple of pseudocircles which forms a triangle is Krupp, then \mathcal{A} is not circularizable.*

We outline the proof: Suppose a realization of \mathcal{A} exists on the sphere. Continuously move the planes spanned by the circles towards the origin. The induced arrangement will eventually become isomorphic to an arrangement of great-circles. Now consider the first event that occurs. As the planes move towards the origin, there is no digon collapse. Since \mathcal{A} is intersecting, no digon is created, and, since all triangles are Krupp, the corresponding intersection points of their planes is already inside S . Therefore, no event can occur – a contradiction.

Proof (first proof of non-circularizability of \mathcal{N}_6^2). The arrangement \mathcal{N}_6^2 is intersecting but not an arrangement of great-pseudocircles (\mathcal{N}_6^2 contains a NonKrupp) and each triangle in \mathcal{N}_6^2 is Krupp. Hence, Theorem 7.16 implies that \mathcal{N}_6^2 is not circularizable. \square

Besides \mathcal{N}_6^2 , there is exactly one other arrangement of 6 pseudocircles (with digons, see Figure 7.10) where Theorem 7.16 implies non-circularizability. For $n = 7$ there are eight arrangements where the theorem applies; but each of them has one of the two $n = 6$ arrangements as a subarrangement.




Figure 7.10: Another intersecting arrangement of 6 pseudocircles (with digons) where Theorem 7.16 applies. The arrangement is minimal non-circularizable and has symmetry 2 (the colors red-orange, blue-green, and purple-black can be exchanged). Triangles are colored gray.

Our second proof of non-circularizability of \mathcal{N}_6^2 is based on Theorem 6.2.

Proof (second proof of non-circularizability of \mathcal{N}_6^2). Suppose that \mathcal{N}_6^2 has a representation as a circle arrangement \mathcal{C} . We refer to circles and intersection points via the labels of the corresponding objects in Figure 7.11(b).

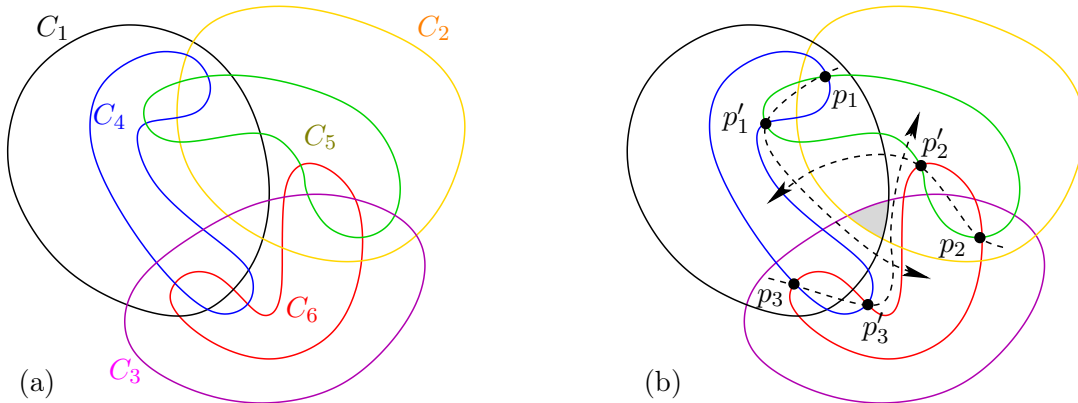


Figure 7.11: (a) The non-circularizable arrangement \mathcal{N}_6^2 . (b) An illustration for the second proof of non-circularizability of \mathcal{N}_6^2 .

Let ℓ_i be the line spanned by p_i and p'_i for $i = 1, 2, 3$. The directed line ℓ_1 intersects C_4 and C_5 in the points p_1, p'_1 and has its second intersection with the yellow circle C_2 between

these points. After p'_1 , the line has to cross C_3 (magenta), C_1 (black), and C_6 (red) in this order, i.e., the line behaves as shown in Figure 7.11. Similarly ℓ_2 and ℓ_3 behave as shown. Let T be the triangle spanned by the intersection points of the three lines ℓ_1, ℓ_2, ℓ_3 . Observe that the gray interior triangle T' of \mathcal{C} is fully contained in T . By applying Theorem 6.2 to the circles C_4, C_5, C_6 , we obtain that ℓ_1, ℓ_2, ℓ_3 meet in a common point, and therefore, T and T' are degenerate. This contradicts the assumption that \mathcal{C} is a realization of \mathcal{N}_6^2 , whence this arrangement is non-circularizable. \square

7.3.3 Non-circularizability of \mathcal{N}_6^3

The arrangement \mathcal{N}_6^3 is shown in Figure 7.2(c) and Figure 7.12(b). To prove its non-circularizability, we again use an incidence lemma. The following lemma is mentioned by Richter-Gebert as a relative of Pappos's Theorem, cf. [RG11, page 26]. Figure 7.12(a) gives an illustration.

Lemma 7.17 ([RG11]). *Let ℓ_1, ℓ_2, ℓ_3 be lines, C'_1, C'_2, C'_3 be circles, and $p_1, p_2, p_3, q_1, q_2, q_3$ be points, such that for $\{i, j, k\} = \{1, 2, 3\}$ point p_i is incident to line ℓ_i , circle C'_j , and circle C'_k , while point q_i is incident to circle C'_i , line ℓ_j , and line ℓ_k . Then C'_1, C'_2 , and C'_3 have a common point of intersection.*

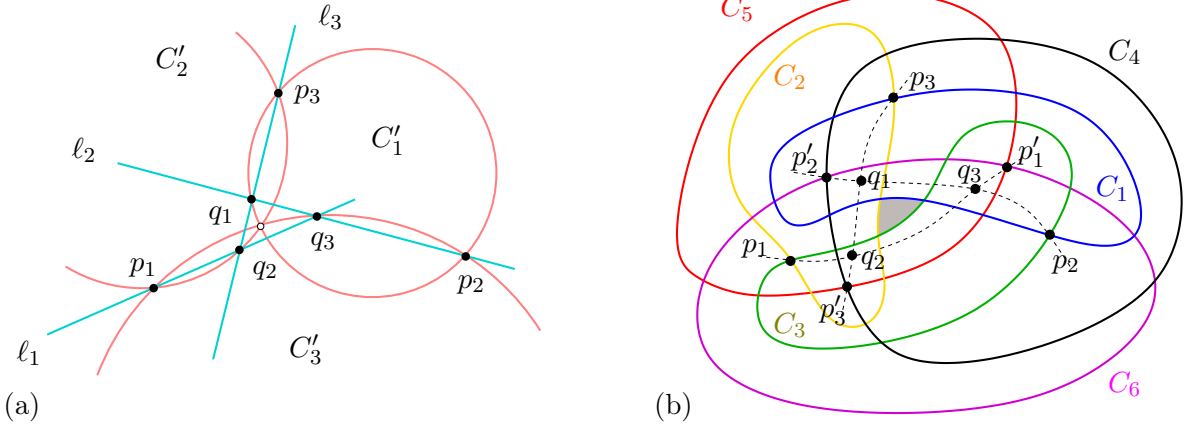


Figure 7.12: (a) An illustration for Lemma 7.17.

(b) The arrangement \mathcal{N}_6^2 with 3 dashed pseudolines illustrating the proof of non-circularizability.

Proof (non-circularizability of \mathcal{N}_6^3). Suppose that \mathcal{N}_6^3 has a representation \mathcal{C} as a circle arrangement in the plane. We refer to circles and intersection points via the label of the corresponding object in Figure 7.12(b). As in the figure, we assume without loss of generality that the triangular cell spanned by C_4, C_5 , and C_6 is the outer cell of the arrangement.

Consider the region $R := R_{24} \cup R_{35}$ where R_{ij} denotes the intersection of the respective interiors of C_i and C_j . The two straight-line segments $p_1p'_1$ and $p_3p'_3$ are fully contained in R_{35} and R_{24} , respectively, and have alternating end points along the boundary of R , hence they cross inside the region $R_{24} \cap R_{35}$.

From rotational symmetry we obtain that the three straight-line segments $p_1p'_1$, $p_2p'_2$, and $p_3p'_3$ intersect pairwise.

For $i = 1, 2, 3$, let ℓ_i denote the line spanned by p_i and p'_i , let q_i denote the intersection-point of ℓ_{i+1} and ℓ_{i+2} , and let C'_i denote the circle spanned by q_i, p_{i+1}, p_{i+2} (indices modulo 3). Note that ℓ_i contains p_i, q_{i+1}, q_{i+2} . These are precisely the conditions for the incidences of points, lines, and circles in Lemma 7.17. Hence, the three circles C'_1, C'_2 , and C'_3 intersect in a common point (cf. Figure 7.12(a)).

Let T be the triangle with corners p_1, p_2, p_3 . Since p_2 and p_3 are on C_1 , and q_1 lies inside of C_1 , we find that the intersection of the interior of C'_1 with T is a subset of the intersection

of the interior of C_1 with T . The respective containments also hold for C'_2 and C_2 and for C'_3 and C_3 . Moreover, since C'_1 , C'_2 , and C'_3 intersect in a common point, the union of the interiors of C'_1 , C'_2 , and C'_3 contains T . Hence, the union of interiors of the C_1 , C_2 , and C_3 also contains T . This shows that in \mathcal{C} there is no face corresponding to the gray triangle; see Figure 7.12(b). This contradicts the assumption that \mathcal{C} is a realization of \mathcal{N}_6^3 , whence the arrangement is non-circularizable. \square

7.4 Additional Arrangements with $n = 6$

In the previous two sections we have exhibited all non-circularizable arrangements with $n \leq 5$ and all non-circularizable intersecting digon-free arrangements with $n = 6$. With automatized procedures we managed to find circle representations of 98% of the connected digon-free arrangements and of 90% of the intersecting arrangements of 6 pseudocircles. Unfortunately, the numbers of remaining candidates for non-circularizability are too large to complete the classification by hand (cf. Section 9.7). In this section we show non-circularizability of a few of the remaining examples which we consider to be interesting. As a criterion for being interesting we used the order of the symmetry group of the arrangement. The symmetry groups have been determined as the automorphism groups of the primal-dual graphs using SageMath [S⁺17a, S⁺17c].

In Chapter 7.4.1 we show non-circularizability of the three intersecting arrangements of $n = 6$ pseudocircles (with digons) depicted in Figure 7.13. The symmetry group of these three arrangements is of order 6. All the remaining examples of intersecting arrangements with $n = 6$, where we do not know about circularizability, have a symmetry group of order at most 3.




Figure 7.13: Three intersecting arrangements of $n = 6$ pseudocircles with symmetry 6. (a) \mathcal{N}_6^{ER} (b) $\mathcal{N}_6^{i6:2}$ (c) $\mathcal{N}_6^{i6:3}$.

In Chapter 7.4.2 we show non-circularizability of the five connected digon-free arrangements of 6 pseudocircles depicted in Figures 7.17 and 7.20. The symmetry group of the three arrangements shown in Figure 7.17 is of order 24 or 8, and the symmetry group of the two arrangements shown in Figure 7.20 is of order 4. All the remaining examples of connected digon-free arrangements with $n = 6$, where we do not know about circularizability, have a symmetry group of order 2 or 1.

7.4.1 Non-circularizability of 3 intersecting arrangements with $n = 6$

In this subsection we prove non-circularizability of the three arrangements \mathcal{N}_6^{ER} , $\mathcal{N}_6^{i6:2}$, and $\mathcal{N}_6^{i6:3}$ shown in Figure 7.13. The non-circularizability of \mathcal{N}_6^{ER} was already shown by Edelsbrunner and Ramos [ER97], the name of the arrangement reflects this fact. The other names are built, such that the subscript of the \mathcal{N} is the number of pseudocircles, the first part of the superscript indicates that the arrangement is intersecting with a symmetry group of order 6, and the number after the colon is the counter. Accordingly, the arrangement \mathcal{N}_6^{ER} can also be denoted as $\mathcal{N}_6^{i6:1}$.

Non-circularizability of the Edelsbrunner–Ramos example \mathcal{N}_6^{ER}

The arrangement \mathcal{N}_6^{ER} is shown in Figure 7.13(a). As in the original proof [ER97] the argument is based on considerations involving angles.

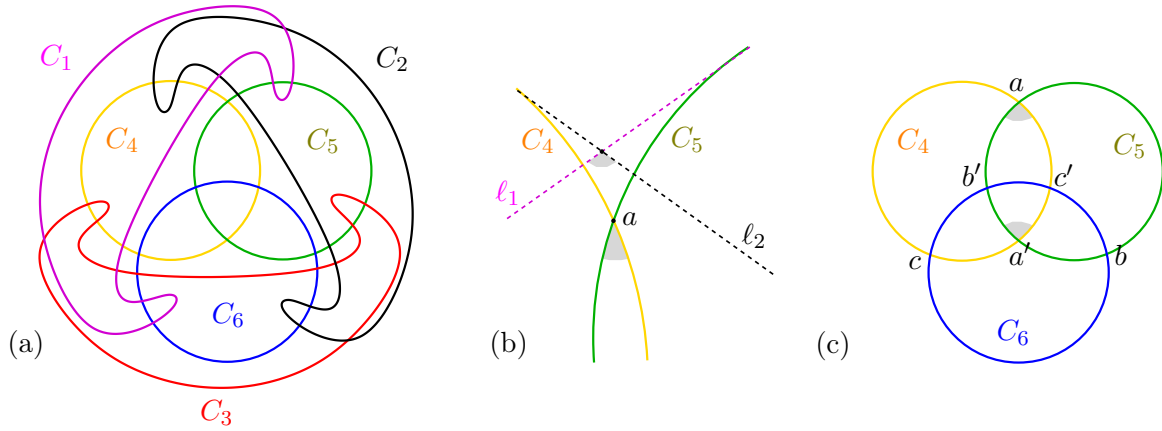


Figure 7.14: (a) The Edelsbrunner–Ramos example \mathcal{N}_6^{ER} .
 (b) Comparing the angle at a and the corresponding angle of T .
 (c) Labels for the vertices of the inner subarrangement \mathcal{A}_I .

Figure 7.14(a) shows a representation of the arrangement \mathcal{N}_6^{ER} consisting of a subarrangement \mathcal{A}_O formed by the three *outer pseudocircles* C_1, C_2, C_3 and a second subarrangement \mathcal{A}_I formed by the three *inner circles* C_4, C_5, C_6 .

Suppose that there is a circle representation \mathcal{C} of \mathcal{N}_6^{ER} . Let \mathcal{C}_O and \mathcal{C}_I be the subarrangements of \mathcal{C} which represent \mathcal{A}_O and \mathcal{A}_I , respectively. For each outer circle C_i from \mathcal{C}_O consider a straight-line segment s_i that connects two points from the two digons which are formed by C_i with inner circles. The segment s_i is fully contained in C_i . Let ℓ_i be the line supporting s_i and let T be the triangle bounded by ℓ_1, ℓ_2 , and ℓ_3 .

We claim that T contains the inner triangle of \mathcal{C}_O . Indeed, if three circles form a NonKrupp where the outer face is a triangle and with each circle we have a line which intersects the two digons incident to the circle, then the three lines form a triangle containing the inner triangular cell of the NonKrupp arrangement.

The inner triangle of \mathcal{C}_O contains the four inner triangles of \mathcal{C}_I . Let a, b, c be the three crossing points on the outer face of the subarrangement \mathcal{C}_I . Comparing the inner angle at a , a crossing of C_4 and C_5 , and the corresponding angle of T , i.e., the angle formed by ℓ_1 and ℓ_2 , we claim that the inner angle at a is smaller. To see this let us assume that the common tangent h of C_4 and C_5 on the side of a is horizontal. Line ℓ_1 has both crossings with C_5 above a and also intersects with C_6 . This implies that the slope of ℓ_1 is positive but smaller than the slope of the tangent at C_5 in a . Similarly, the slope of ℓ_2 is negative but larger than the slope of the tangent at C_4 in a . This implies the claim, see Figure 7.14(b).

The respective statements hold for the inner angles at b and c , and the corresponding angles of T . Since the sum of angles of T is π , we conclude that the sum of the inner angles at a, b , and c is less than π .

The sum of the inner angles at a, b, c equals the sum of inner angles at a', b', c' , see Figure 7.14(c). This sum, however, clearly exceeds π . The contradiction shows that \mathcal{N}_6^{ER} is not circularizable.

Non-circularizability of $\mathcal{N}_6^{i6:2}$

The arrangement $\mathcal{N}_6^{i6:2}$ is shown in Figure 7.13(b) and again in Figure 7.15(a). This figure also shows some shaded triangles, three of them are gray and three are pink.




Figure 7.15: (a) The arrangement $\mathcal{N}_6^{i6:2}$ with some triangle faces emphasized.
 (b) After collapsing the shaded triangles.
 (c) After moving the point p_1 to infinity.

Suppose that $\mathcal{N}_6^{i6:2}$ has a circle representation \mathcal{C} . Each of C_1 , C_2 , and C_3 has two triangles and no digon on its interior boundary. One of the two triangles is gray the other pink. Lemma 7.15 allows to shrink the three circles C_1 , C_2 , C_3 of \mathcal{C} into their respective interiors such that in each case the shrinking makes a pink triangle collapse. Let p_i be the point of triple intersection of C_i , for $i = 1, 2, 3$. Further shrinking C_i towards p_i makes another triangle collapse. At this second collapse two triangles disappear, one of them a gray one, and C_i gets incident to p_{i-1} (with indices modulo 3). Having done this for each of the three circles C_1, C_2, C_3 yields a circle representation for the (non-simple) arrangement shown in Figure 7.15(b).

To see that this arrangement has no circle representation apply a Möbius transformation that maps the point p_1 to the point ∞ of the extended complex plane. This transforms the four circles C_1 , C_2 , C_4 , C_5 , which are incident to p_1 , into lines. The two remaining circles C_3 and C_6 intersect in p_2 and p_3 . The lines of C_2 and C_5 both have their second intersections with C_3 and C_6 separated by p_2 , hence, they both avoid the lens formed by C_3 and C_6 . The line of C_1 has its intersections with C_2 and C_5 in the two components of the gray double-wedge of C_2 and C_5 , see Figure 7.15(c). Therefore, the slope of C_1 belongs to the slopes of the double-wedge. However, the line of C_1 has its second intersections with C_3 and C_6 on the same side of p_3 and, therefore, it has a slope between the tangents of C_3 and C_6 at p_3 . These slopes do not belong to the slopes of the gray double-wedge. This contradiction shows that a circle representation \mathcal{C} of $\mathcal{N}_6^{i6:2}$ does not exist.

Non-circularizability of $\mathcal{N}_6^{i6:3}$

The arrangement $\mathcal{N}_6^{i6:3}$ is shown in Figure 7.13(c) and again in Figure 7.16(a). This figure also shows some shaded triangles, three of them are gray and three are pink.

Suppose that $\mathcal{N}_6^{i6:2}$ has a circle representation \mathcal{C} . Each of C_1 , C_2 and C_3 has two triangles and no digon on its exterior boundary (excluding the outer face). One of the two triangles is gray the other is pink (we disregard the exterior triangle because it will not appear in a first flip when expanding a circle). The dual form of Lemma 7.15 allows to expand the three circles C_1 , C_2 , C_3 of \mathcal{C} into their respective exteriors such that in each case the expansion makes a pink triangle collapse.

Figure 7.16(b) shows a pseudocircle representation of the arrangement after this first phase of collapses. In the second phase, we modify the circles C_i , for $i = 4, 5, 6$. We explain what happens to C_5 , the two other circles are treated the same with respect to the rotational symmetry. Consider the circle C'_5 , which contains p_1 and shares p_3 and the red point with C_5 . This circle is obtained by shrinking C_5 on one side of the line containing p_3 and the red point, and by




Figure 7.16: (a) The arrangement $\mathcal{N}_6^{i6:3}$ with some triangle faces emphasized.
 (b) After collapsing the pink shaded triangles.
 (c) A detail of the arrangement after the second phase of collapses.

expanding C_5 on the other side of the line. It is easily verified that the collapse of the gray triangle at p_1 is the first event in this process.

Figure 7.16(c) shows the inner triangle formed by C_1 , C_2 , and C_3 together with parts of C'_4 , C'_5 , and C'_6 . At each of the three points, the highlighted red angle is smaller than the highlighted gray angle. However, the red angle at p_i is formed by the same two circles as the gray angle at p_{i+1} , whence, the two angles are equal. This yields a contradictory cyclic chain of inequalities. The contradiction shows that a circle representation \mathcal{C} of $\mathcal{N}_6^{i6:3}$ does not exist.

7.4.2 Non-circularizability of 5 connected digon-free arrangements with $n = 6$

We first prove non-circularizability of the three arrangements \mathcal{N}_6^{c24} , $\mathcal{N}_6^{c8:1}$, and $\mathcal{N}_6^{c8:2}$ shown in Figure 7.17.




Figure 7.17: A digon-free connected arrangement of $n = 6$ pseudocircles with symmetry group of order 24, and two with a symmetry group of order 8: (a) \mathcal{N}_6^{c24} (b) $\mathcal{N}_6^{c8:1}$ (c) $\mathcal{N}_6^{c8:2}$.

Non-circularizability of \mathcal{N}_6^{c24} and $\mathcal{N}_6^{c8:1}$

The proof of non-circularizability of the two arrangements is based on Miquel's Theorem. For proofs of the theorem we refer to [RG11].

Theorem 7.18 (Miquel's Theorem). *Let C_1, C_2, C_3, C_4 be four circles and let $C_1 \cap C_2 = \{a, w\}$, $C_2 \cap C_3 = \{b, x\}$, $C_3 \cap C_4 = \{c, y\}$, and $C_4 \cap C_1 = \{d, z\}$. If there is a circle C containing a, b, c, d , then there is a circle C' containing w, x, y, z .*

The arrangement \mathcal{N}_6^{c24} is shown in Figure 7.17(a) and again in Figure 7.18(a). Suppose that \mathcal{N}_6^{c24} has a circle representation \mathcal{C} . Circle C_5 has six triangles in its exterior. These triangles are all incident to either the crossing of C_1 and C_4 or the crossing of C_2 and C_3 . Hence, by Lemma 7.15 we can grow C_5 into its exterior to get two triple intersection points




Figure 7.18: (a) The arrangement \mathcal{N}_6^{c24} .
 (b) \mathcal{N}_6^{c24} after collapsing four triangles.
 (c) The arrangement $\mathcal{N}_6^{c8:1}$.

$p_1 = C_1 \cap C_4 \cap C_5$ and $p_2 = C_2 \cap C_3 \cap C_5$. The situation in the interior of C_6 is identical to the situation in the exterior of C_5 . Hence, by shrinking C_6 we get two additional triple intersection points $p_3 = C_1 \cap C_2 \cap C_6$ and $p_4 = C_3 \cap C_4 \cap C_6$. This yields the (non-simple) arrangement shown in Figure 7.18(b). Now grow the circles C_1, C_2, C_3, C_4 to the outside while keeping each of them incident to its two points p_i , this makes them shrink into their inside at the ‘short arc’. Upon this growth process, the gray crossings p_{12}, p_{23}, p_{34} , and p_{14} move away from the blue circle C_6 . Hence, the process can be continued until the upper and the lower triangles collapse, i.e., until p_{12} and p_{34} both are incident to C_5 . Note that we do not care about p_{23} and p_{14} , they may have passed to the other side of C_5 . The collapse of the upper and the lower triangle yields two additional triple intersection points $C_1 \cap C_2 \cap C_5$ and $C_3 \cap C_4 \cap C_5$. The circles C_1, C_2, C_3, C_4 together with C_5 in the role of C form an instance of Miquel’s Theorem (Theorem 7.18). Hence, there is a circle C' traversing the four points p_3, p_{14}, p_4, p_{23} . The points p_3, p_4 partition C' into two arcs, each containing one of the gray points p_{14}, p_{23} . Now C' shares the points p_3 and p_4 with C_6 while the gray points p_{14}, p_{23} are outside of C_6 . This is impossible, whence there is no circle representation of \mathcal{N}_6^{c24} .

The arrangement $\mathcal{N}_6^{c8:1}$ is shown in Figure 7.17(b) and again in Figure 7.18(c). The proof of non-circularizability of this arrangement is exactly as the previous proof, just replace \mathcal{N}_6^{c24} by $\mathcal{N}_6^{c8:1}$ and think of an analog of Figure 7.18(b). \square

Originally, we were aiming at deriving the non-circularizability of \mathcal{N}_6^{c24} as a corollary to the following theorem. Turning things around we now prove it as a corollary to the non-circularizability of \mathcal{N}_6^{c24} . We say that a polytope P has the combinatorics of the cube if P and the cube have isomorphic face lattices. The graph of the cube is bipartite, hence, we can speak of the white and black vertices of a polytope with the combinatorics of the cube.

Theorem 7.19. *Let S be a sphere. There is no polytope P with the combinatorics of the cube such that the black vertices of P are inside S and the white vertices of P are outside S .*

Proof. Suppose that there is such a polytope P . Let \mathcal{E} be the arrangement of planes spanned by the six faces of P and let \mathcal{C} be the arrangement of circles obtained from the intersection of \mathcal{E} and S . This arrangement is isomorphic to \mathcal{N}_6^{c24} . To see this, consider the eight triangles of \mathcal{N}_6^{c24} corresponding to the black and gray points of Figure 7.18(b). Triangles corresponding to black points are Krupp and triangles corresponding to gray points are NonKrupp. By Fact 6.1 this translates to corners of P being outside and inside S , respectively. \square

Non-circularizability of $\mathcal{N}_6^{c8:2}$

The arrangement $\mathcal{N}_6^{c8:2}$ is shown in Figure 7.17(c) and again in Figure 7.19(a).




Figure 7.19: (a) The arrangement $\mathcal{N}_6^{c8:2}$ with four gray triangles.
 (b) $\mathcal{N}_6^{c8:2}$ after collapsing the gray triangles.
 (c) After moving the point p to infinity.

Suppose that $\mathcal{N}_6^{c8:2}$ has a circle representation \mathcal{C} . Circle C_1 only has two triangles in the exterior, in the figure they are gray. Circle C_6 only has two triangles in the interior. With Lemma 7.15 these four triangles can be collapsed into points of triple intersection. This results in a (non-simple) arrangement as shown in Figure 7.19(b). Note that we do not care, whether the circles C_1 and C_6 cross or not.

Apply a Möbius transformation that maps the point $p = C_4 \cap C_5 \cap C_6$ to the point ∞ of the extended complex plane. This maps C_4 , C_5 , and C_6 to lines, while C_1 , C_2 , and C_3 are mapped to circles. From the order of crossings, it follows that the situation is essentially as shown in Figure 7.19(c). This Figure also shows the line ℓ_1 through the two intersection points of C_1 and C_2 , the line ℓ_2 through the two intersection points of C_2 and C_3 , and the line ℓ_3 through the two intersection points of C_3 and C_1 . The intersection points of ℓ_3 with the other two are separated by the two defining points of ℓ_3 . According to Theorem 6.2, however, the three lines should share a common point. The contradiction shows that there is no circle representation of $\mathcal{N}_6^{c8:2}$.

Non-circularizability of $\mathcal{N}_6^{c4:1}$ and $\mathcal{N}_6^{c4:2}$

The arrangements $\mathcal{N}_6^{c4:1}$ and $\mathcal{N}_6^{c4:2}$ are shown in Figure 7.20 and again in Figure 7.21. These are the only two connected digon-free arrangements of 6 pseudocircles with a symmetry group of order 4 which are not circularizable. The two proofs of non-circularizability are very similar.




Figure 7.20: Two non-circularizable arrangements of $n = 6$ pseudocircles with a symmetry group of order 4. The arrangements are denoted as (a) $\mathcal{N}_6^{c4:1}$ and (b) $\mathcal{N}_6^{c4:2}$.

Suppose that $\mathcal{N}_6^{c4:1}$ has a circle representation \mathcal{C} . In Figure 7.21(a) the pseudocircles C_5 and C_6 each have two gray triangles on the outside and these are the only triangles on the outside of the two pseudocircles. With Lemma 7.15 the respective circles in \mathcal{C} can be grown until the gray triangles collapse into points of triple intersection or until a digon flip occurs. In the case of a digon flip C_5 and C_6 become intersecting, and no further triangles incident to C_5 and C_6

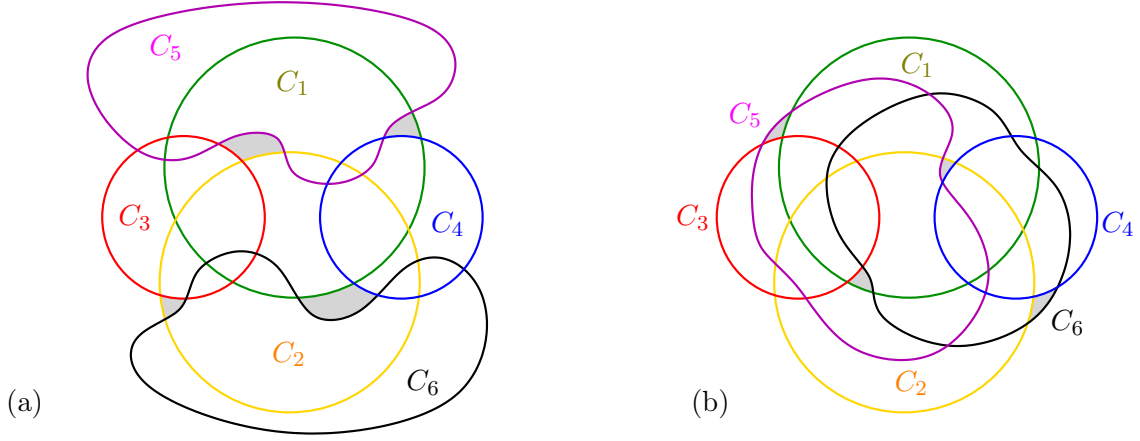


Figure 7.21: (a) The arrangement $\mathcal{N}_6^{c4:1}$ with four gray triangles.
(b) The arrangement $\mathcal{N}_6^{c4:2}$ with four gray triangles.

are created. Therefore, it is possible to continue the growing process until the four triangles collapse. In the following, we do not care whether the digon flip occurred during the growth process, i.e., whether C_5 and C_6 intersect. The four points of triple intersection are the points p_1, p_2, p_3 , and p_4 in Figure 7.22(a).

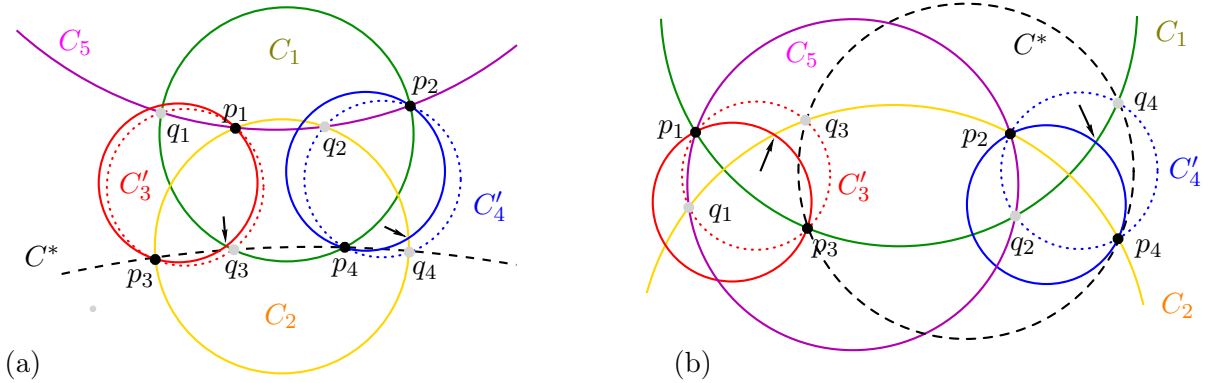


Figure 7.22: (a) Illustration of the non-circularizability proof for $\mathcal{N}_6^{c4:1}$.
(b) Illustration of the non-circularizability proof for $\mathcal{N}_6^{c4:2}$.

There is a circle C'_3 which shares the points p_1 and p_3 with C_3 and also contains the point q_1 , which is defined as the intersection point of C_1 and C_5 inside C_3 . Similarly there is a circle C'_4 which shares the points p_2 and p_4 with C_4 and also contains the point q_2 , which is defined as the intersection point of C_2 and C_5 inside C_4 . By construction C_5 is incident to one intersection point of each of the pairs C_1, C'_3 , and C'_3, C_2 , and C_2, C'_4 and C'_4, C_1 . Miquel's Theorem (Theorem 7.18) implies that there is a circle C^* through the second intersection points of these pairs. It can be argued that on C^* the two points p_3 and p_4 separate q_3 and q_4 . The circle C_6 shares the points p_3 and p_4 with C^* and contains the crossing of the pairs C_1, C_3 and C_2, C_4 which are 'close to' q_3 and q_4 in its interior (the two points are emphasized by the arrows in the figure). Hence p_3 and p_4 are separated by q_3 and q_4 . This is impossible, whence $\mathcal{N}_6^{c4:1}$ is not circularizable.

Suppose that $\mathcal{N}_6^{c4:2}$ has a circle representation \mathcal{C} . In Figure 7.21(b) the pseudocircles C_3 and C_4 each have two gray triangles on the outside and these are the only triangles on the outside of the two pseudocircles. By Lemma 7.15 the respective circles in \mathcal{C} can be grown to make the gray triangles collapse into points of triple intersection (we do not care whether during the growth

process C_3 and C_4 become intersecting). The four points of triple intersection are the points p_1 , p_2 , p_3 , and p_4 in Figure 7.22(b).

There is a circle C'_3 which shares the points p_1 and p_3 with C_3 and also contains the point q_1 , this point q_1 is the intersection point of C_2 and C_5 inside C_3 . Similarly there is a circle C'_4 which shares the points p_2 and p_4 with C_4 and also contains the point q_2 , this point q_2 is the intersection point of C_1 and C_5 inside C_4 . By construction C_5 contains one intersection point of the pairs C_1, C'_3 , and C'_3, C_2 , and C_2, C'_4 and C'_4, C_1 . Miquel's Theorem (Theorem 7.18) implies that there is a circle C^* through the second intersection points of these pairs. It can be argued that on C^* the points q_3 and q_4 belong to the same of the two arcs defined by the pair p_3, p_4 . The circle C_6 shares the points p_3 and p_4 with C^* and has the crossing of the pair C_2, C_3 which is outside C^* in its inside and the crossing of the pair C_1, C_4 which is inside C^* in its outside (the two points are emphasized by the arrows in the figure). This is impossible, whence $\mathcal{N}_6^{c4:2}$ is not circularizable.

7.5 Enumeration and Asymptotics

Recall from Chapter 5, that the *primal graph* of a connected arrangement of $n \geq 2$ pseudocircles is the plane graph whose vertices are the crossings of the arrangement and edges are pseudoarcs, i.e., pieces of pseudocircles between consecutive crossings. The primal graph of an arrangement is a simple graph if and only if the arrangement is digon-free.

The *dual graph* of a connected arrangement of $n \geq 2$ pseudocircles is the dual of the primal graph, i.e., vertices correspond to the faces of the arrangement and edges correspond to pairs of faces which are adjacent along a pseudoarc. The dual graph of an arrangement is simple if and only if the arrangement remains connected after the removal of any pseudocircle. In particular, dual graphs of intersecting arrangements are simple.

The *primal-dual graph* of a connected arrangement of $n \geq 2$ pseudocircles has three types of vertices; the vertices correspond to crossings, pseudoarcs, and faces of the arrangement. Edges represent incident pairs of a pseudoarc and a crossing, and of a pseudoarc and a face¹.

Figure 7.23 shows the primal graph, the dual graph, and the primal-dual graph of the NonKrupp arrangement.

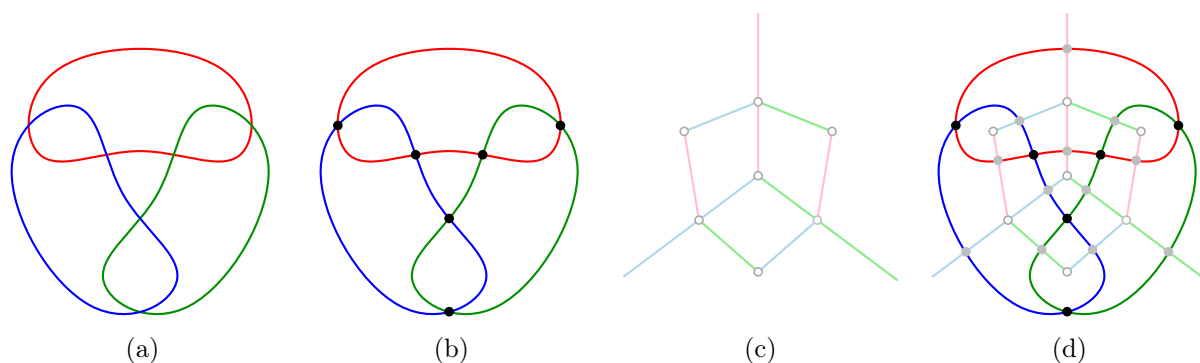


Figure 7.23: (a) An arrangement and its (b) primal graph, (c) dual graph, and (d) primal-dual graph.

We now show that crossing free embeddings of dual graphs and primal-dual graphs on the sphere are unique. This will allow us to disregard the embedding and work with abstract graphs.

If G is a subdivision of a 3-connected graph H , then we call G *almost 3-connected*. If H is planar and, hence, has a unique embedding on the sphere, then the same is true for G .

¹In retrospect, we believe that adding the face-crossing edges to the primal-dual graph would have simplified the approach with respect to theory as well as to implementation.

In the dual and the primal-dual graph of an arrangement of pseudocircles the only possible 2-separators are the two neighboring vertices of a vertex corresponding to a digon. It follows that these graphs are almost 3-connected. We conclude the following.

Proposition 7.20. *The dual graph of a simple intersecting arrangement of $n \geq 2$ pseudocircles has a unique embedding on the sphere.*

Proposition 7.21. *The primal-dual graph of a simple connected arrangement of $n \geq 2$ pseudocircles has a unique embedding on the sphere.*

Note that the statement of Proposition 7.20 clearly holds for $n = 2$, where the dual graph is the 4-cycle.

7.5.1 Enumeration of Arrangements

The database of all intersecting arrangements of up to $n = 7$ pseudocircles was generated with a recursive procedure. Arrangements of pseudocircles were represented by their dual graphs. The recursion was initiated with the unique arrangement of two intersecting pseudocircles. Given the dual of an arrangement we used a procedure which generates all possible extensions by one additional pseudocircle. The procedure is based on the observation that a pseudocircle in an intersecting arrangement of n pseudocircles corresponds to a cycle of length $2n - 2$ in the dual graph. A problem is that an arrangement of n pseudocircles is generated up to n times. Since the embedding of the dual graph is unique, we could use the canonical labeling provided by the Graph-package of SageMath [S⁺17a] to check whether an arrangement was found before².

Another way for obtaining a database of all intersecting arrangements of n pseudocircles for a fixed value of n , is to start with an arbitrary intersecting arrangement of n pseudocircles and then perform a recursive search in the flip-graph using the triangle flip operation (cf. Chapter 7.5.5).

Recall that the dual graph of a connected arrangement contains multiple edges if the removal of one of the pseudocircles disconnects the arrangement. Hence, to avoid problems with non-unique embeddings, we modeled connected arrangements with their primal-dual graphs. To generate the database of all connected arrangements for $n \leq 6$, we used the fact that the flip-graph is connected, when both triangle flips and digon flips are used (cf. Chapter 6.3). The arrangements were created with a recursive search on the flip-graph.

7.5.2 Generating Circle Representations

Having generated the database of arrangements of pseudocircles, we were then interested in identifying the circularizable and the non-circularizable ones.

Our first approach was to generate arrangements of circles C_1, \dots, C_n with centers (x_i, y_i) and radius r_i by choosing triples x_i, y_i, r_i at random from $\{1, \dots, K\}$ for a fixed constant $K \in \mathbb{N}$. In the database the entries corresponding to the generated arrangements were marked circularizable. Later we used known circle representations to find new ones by perturbing values. In particular, whenever a new circle arrangement was found, we tried to locally modify the parameters to obtain further arrangements. With these quantitative approaches we managed to break down the list for $n = 5$ to few “hard” examples, which were then treated “by hand”.

For later computations on $n = 6$ (and $n = 7$), we also used the information from the flip-graph on all arrangements of pseudocircles. In particular, to find realizations for a “not-yet-realized” arrangement, we used neighboring arrangements which had already been realized for perturbations. This approach significantly improved the speed of realization.

Another technique to speedup our computations was to use floating point operations and, whenever a solution suggested that an additional arrangement is circularizable, we verified the

²We recommend the Sage Reference Manual on Graph Theory [S⁺17c] and its collection of excellent examples.

solution using exact arithmetics. Note that the intersection points of circles, described by integer coordinates and integer radii, have algebraic coordinates, and can therefore be represented by minimal polynomials. All computations were done using the computer algebra system SageMath [S⁺17a]³.

As some numbers got quite large during the computations, we took efforts to reduce the “size” of the circle representations, i.e., the maximum over all parameters $|x_i|, |y_i|, r_i$. It turned out to be effective to scale circle arrangements by a large constant, perturb the parameters, and divide all values by the greatest common divisor. This procedure allowed to reduce the number of bits significantly when storing the circle $(x - a)^2 + (y - b)^2 = r^2$ with $a, b, r \in \mathbb{Z}$.

7.5.3 The Circularizability Problem as System of Polynomial Inequalities

We here describe, for a given arrangement of pseudocircles \mathcal{A} in the Euclidean plane, a system of polynomial inequalities over real-valued variables which has a solution if and only if the arrangement is circularizable. This model not only shows the containment in the complexity class ETR (cf. Corollary 7.7), one might also use algebraic or numeric solvers to find a solution of this system, and hence a realization of \mathcal{A} . In theory, one could also prove non-circularizability based on this model, in practice, however, no solver/algorithm (known to us) can be used to prove infeasibility of such an instance in reasonable time.

In our system of inequalities, each circle C_i in an arrangement \mathcal{C} of circles in the Euclidean plane is represented by its center $c_i = (x_i, y_i) \in \mathbb{R}^2$ and its radius $r_i > 0$. Since each two circles C_i, C_j are distinct, we clearly have

$$\|c_i - c_j\|^2 + (r_i - r_j)^2 > 0, \quad (7.1)$$

where $\|\cdot\|^2$ denotes the squared Euclidean norm, that is, a quadratic polynomial in the respective x - and y -coordinates. Pause to note that the left-hand-side is clearly non-negative as the sum over quadratic terms and that, given $X \geq 0$, the strict inequality $X > 0$ is equivalent to $X - \varepsilon^2 \geq 0 \wedge \varepsilon \cdot \delta = 1$, where ε and δ denote auxiliary variables. This transformation can be applied to any strict inequation and allows to rewrite any system of inequalities only using “ \leq ”-symbols.

Two circles C_i and C_j intersect if and only if

$$(r_i - r_j)^2 \leq \|c_i - c_j\|^2 \leq (r_i + r_j)^2, \quad (7.2)$$

while one of the two equalities holds if and only if the two circles touch. Since our focus is on simple arrangements, we can forbid touching circles using the fact that $X \neq Y$ is equivalent to $X - Y = \varepsilon \wedge \varepsilon \cdot \delta = 1$, where ε and δ denote auxiliary variables. In particular, we write:

$$(r_i - r_j)^2 - \|c_i - c_j\|^2 = \varepsilon_{ij} \quad \text{and} \quad \varepsilon_{ij} \cdot \delta_{ij} = 1, \quad (7.3)$$

$$(r_i + r_j)^2 - \|c_i - c_j\|^2 = \varepsilon'_{ij} \quad \text{and} \quad \varepsilon'_{ij} \cdot \delta'_{ij} = 1. \quad (7.4)$$

From the given arrangement we know which pairs of circles intersect and which do not. For a pair C_i, C_j of intersecting circles we denote the two points of intersections by $p_{ij} = (x_{ij}, y_{ij})$ and $p_{ji} = (x_{ji}, y_{ji})$. When orienting the circles counter-clockwise, the point p_{ij} can be considered as the point, where C_i enters the disk bounded by C_j , and vice versa. A point p which lies on a circle C_i fulfills

$$\|c_i - p\|^2 = r_i^2. \quad (7.5)$$

For a pair of circles C_i, C_j that intersect, the inequalities from (7.2) can be omitted as the existence of the points p_{ij}, p_{ji} asserts an intersection.

³For more details, we refer to the Sage Reference Manual on Algebraic Numbers and Number Fields [S⁺17b].

For a pair of circles C_i, C_j that do not intersect, however, the negation of the inequalities (7.2) are indeed necessary to prevent unintended intersections. Thus, we use the auxiliary variables $\varepsilon_{ij}, \varepsilon'_{ij}$ from inequalities (7.3) and (7.4) to assert

$$\varepsilon_{ij} \cdot \varepsilon'_{ij} > 0. \quad (7.6)$$

For a fixed circle C_i , the order of the points of intersections along C_i are prescribed by the arrangement and all these points are in convex position. For each three points of intersection p_{ij}, p_{ik}, p_{il} on C_i , the sign of the determinant of the 3×3 matrix

$$\begin{pmatrix} 1 & 1 & 1 \\ x_{ij} & x_{ik} & x_{il} \\ y_{ij} & y_{ik} & y_{il} \end{pmatrix}$$

is prescribed by the arrangement. Note that this determinant is a quadratic polynomial in the respective x - and y -coordinates. Since all points lie on the circle C , no three lie on a common line and, therefore, the sign of the determinant is either strictly positive or strictly negative.

Altogether, we obtain a system of quadratic inequalities of polynomial size which has a solution if and only if the arrangement \mathcal{A} is circularizable.

7.5.4 Counting Arrangements

Projective arrangements of pseudolines are also known as projective abstract order types or oriented matroids of rank 3. The precise numbers of such arrangements are known for $n \leq 11$ [Knu92, AAK02, Kra03, AK06] (cf. sequences A006248 and A018242 on the OEIS [Slo]). Hence the numbers of great-pseudocircle arrangements given in Table 7.1 are not new. Moreover, it is well-known that there are $2^{\Theta(n^2)}$ arrangements of pseudolines and only $2^{\Theta(n \log n)}$ arrangements of lines [GP93, FG18]. Those bounds directly translate to arrangements of great-pseudocircles, and, in particular, there are at least $2^{\Omega(n^2)}$ and $2^{\Omega(n \log n)}$ arrangements of pseudocircles and circles, respectively. In this subsection we show that the number of such arrangements is also bounded from above by $2^{O(n^2)}$ and $2^{O(n \log n)}$, respectively.

Proposition 7.22. *There are $2^{\Theta(n^2)}$ arrangements of n pseudocircles.*

Proof. The primal-dual graph of a connected arrangement of n pseudocircles is a plane quadrangulation on $O(n^2)$ vertices. A quadrangulation can be extended to a triangulation by inserting a diagonal edge in every quadrangular face. It is well-known that the number of triangulations on s vertices is $2^{\Theta(s)}$ [Tut62]. Hence, the number of connected arrangements of n pseudocircles is bounded by $2^{O(n^2)}$.

Now observe that every not necessarily connected arrangement on n pseudocircles can be extended by n additional pseudocircles to a connected arrangement on $2n$ pseudocircles. Since each connected arrangement on $2n$ pseudocircles yields at most $\binom{2n}{n}$ distinct subarrangement of size n , there are at most $2^{O(n^2)}$ arrangements on n pseudocircles. \square

Proposition 7.23. *There are $2^{\Theta(n \log n)}$ arrangements of n circles.*

The proof relies on a bound for the number of cells in an arrangement of zero sets of polynomials (the underlying theorem is associated with the names Oleinik-Petrovsky, Milnor, Thom, and Warren [PO49, Mil64, Tho65, War68]). The argument is similar to the one given by Goodman and Pollack [GP86] to bound the number of arrangements of lines, see also Chapter 6.2 of Matoušek's book [Mat02].

Proof. An arrangement \mathcal{C} of n circles on the unit sphere \mathbb{S} is induced by the intersection of n planes $\mathcal{E} = \{E_1, \dots, E_n\}$ in 3-space with \mathbb{S} . Plane E_i can be described by the linear equation $a_i x + b_i y + c_i z + d_i = 0$ for some reals a_i, b_i, c_i, d_i ; we call them the parameters of E_i . Below we define a polynomial P_{ijk} of degree 6 in the parameters of the planes, such that $P_{ijk} = 0$ iff the three circles $E_i \cap \mathbb{S}$, $E_j \cap \mathbb{S}$, and $E_k \cap \mathbb{S}$ have a common point of intersection. We also define a polynomial Q_{ij} of degree 8 in the parameters of the planes, such that $Q_{ij} = 0$ iff the circles $E_i \cap \mathbb{S}$ and $E_j \cap \mathbb{S}$ touch.

Transforming an arrangement \mathcal{C} into an arrangement \mathcal{C}' in a continuous way corresponds to a curve γ in \mathbb{R}^{4n} from the parameter vector of \mathcal{E} to the parameter vector of \mathcal{E}' . If a triangle flip or a digon flip occurs when transforming \mathcal{C} to \mathcal{C}' , then γ intersects the zero set of a polynomial P_{ijk} or Q_{ij} . Hence, all the points in a fixed cell of the arrangement defined by the zero set of the polynomials P_{ijk} or Q_{ij} with $1 \leq i < j < k \leq n$ are parameter vectors of isomorphic arrangements of circles.

The number of cells in \mathbb{R}^d induced by the zero sets of m polynomials of degree at most D is upper bounded by $(50Dm/d)^d$ (Theorem 6.2.1 in [Mat02]). Consequently the number of non-isomorphic arrangements of n circles, is bounded by $(50 \cdot 8 \cdot 2 \binom{n}{3} / 4n)^{4n}$ which is $n^{O(n)}$.

For the definition of the polynomial P_{ijk} we first note (see e.g. [RG11, Section 12.3]) that the homogeneous coordinates of the point I_{ijk} of intersection of the three planes E_i, E_j, E_k are given by⁴

$$\bigotimes_4((a_i, b_i, c_i, d_i), (a_j, b_j, c_j, d_j), (a_k, b_k, c_k, d_k)),$$

where \bigotimes_n denotes the $(n-1)$ -ary analogue of the cross product in \mathbb{R}^n

$$\bigotimes_n(\mathbf{v}_1, \dots, \mathbf{v}_{n-1}) := \begin{vmatrix} v_1^{(1)} & \dots & v_{n-1}^{(1)} & \mathbf{e}_1 \\ \vdots & \ddots & \vdots & \vdots \\ v_1^{(n)} & \dots & v_{n-1}^{(n)} & \mathbf{e}_n \end{vmatrix}.$$

For more information on the $(n-1)$ -ary analogue of the cross product, we refer e.g. to [Wikb].

Each component of I_{ijk} is a cubic polynomial in the parameters of the three planes. Since a point in 3-space, described homogeneous coordinates (x, y, w, λ) , lies on the unit sphere \mathbb{S} if and only if $x^2 + y^2 + z^2 - \lambda^2 = 0$, we get a polynomial P_{ijk} of degree 6 in the parameters of the planes such that $P_{ijk} = 0$ iff $I_{ijk} \in \mathbb{S}$.

To define the polynomial Q_{ij} we need some geometric considerations. Note that the two circles $E_i \cap \mathbb{S}$ and $E_j \cap \mathbb{S}$ touch if and only if the line $L_{ij} = E_i \cap E_j$ is tangential to \mathbb{S} . Let E_{ij}^* be the plane normal to L_{ij} which contains the origin. The point I_{ij}^* of intersection of the three planes E_i, E_j, E_{ij}^* is on \mathbb{S} if and only if $E_i \cap \mathbb{S}$ and $E_j \cap \mathbb{S}$ touch.

A vector N_{ij} which is parallel to L_{ij} can be obtained as $\bigotimes_3((a_i, b_i, c_i), (a_j, b_j, c_j))$. The components of N_{ij} are polynomials of degree 2 in the parameters of the planes in \mathcal{E} . The components of N_{ij} are the first three parameters of E_{ij}^* ; the fourth parameter is zero. The homogeneous components of I_{ij}^* are obtained by a \bigotimes_4 from the parameters of the three planes E_i, E_j , and E_{ij}^* . Since the parameters of E_{ij}^* are polynomials of degree 2, the components of I_{ij}^* are polynomials of degree 4. Finally we have to test whether $I_{ij}^* \in \mathbb{S}$, this makes Q_{ij} a polynomial of degree 8 in the parameters of the planes. \square

7.5.5 Connectivity of the Flip-Graph

Given two arrangements of n circles, one can continuously transform one arrangement into the other. During this transformation (the combinatorics of) the arrangement changes whenever a triangle flip or a digon flip occurs.

⁴If the three planes E_i, E_j, E_k intersect in a common line, we still take the expression as a definition for I_{ijk} , i.e., the homogeneous coordinates are all zero.

Snoeyink and Hershberger [SH91] showed an analog for arrangements of pseudocircles: Given two arrangements on n pseudocircles, a sequence of digon flips and triangles flips can be applied to transform one arrangement into the other. In other words, they have proven connectivity of the *flip-graph* of arrangements of n pseudocircles, which has all non-isomorphic arrangements as vertices and edges between arrangements that can be transformed into each other using only a single flip operation.

For arrangements of (pseudo)lines, it is well-known that the triangle flip-graph is connected (see e.g. [FW00]). A triangle flip in an arrangement of (pseudo)lines corresponds to an operation in the corresponding arrangement of great-(pseudo)circles where two “opposite” triangles are flipped simultaneously.

With an idea as in the proof of the Great-Circle Theorem (Theorem 7.1) we can show that the flip-graph for arrangements of circles is connected.

Theorem 7.24. *The triangle flip-graph on the set of all intersecting (digon-free) arrangements of n circles is connected for every $n \in \mathbb{N}$.*

Proof. Consider an intersecting arrangement of circles \mathcal{C} on the unit sphere \mathbb{S} . Imagine the planes of $\mathcal{E}(\mathcal{C})$ moving towards the origin. To be precise, for time $t \geq 1$ let $\mathcal{E}_t := \{1/t \cdot E : E \in \mathcal{E}(\mathcal{C})\}$. During this process only triangle flips occur as the arrangement is already intersecting and eventually the point of intersection of any three planes of \mathcal{E}_t is in the interior of the unit sphere \mathbb{S} . Thus, in the circle arrangement obtained by intersecting the moving planes \mathcal{E}_t with \mathbb{S} every triple of circles forms a Krupp, that is, the arrangement becomes a great-circle arrangement. Since the triangle flip-graph of line arrangements is connected, we can use triangle flips to get to any other great-circle arrangement. Note that, due to Fact 6.2, no digon occurs in the arrangement during the whole process in the setting of digon-free arrangements.

Consequently, any two arrangements of circles \mathcal{C} and \mathcal{C}' can be flipped to the same great-circle arrangement without digons to occur, and the statement follows. \square

Based on the computational evidence for $n \leq 7$, we conjecture that the following is true.

Conjecture 7.2. *The triangle flip-graph on the set of all intersecting (digon-free) arrangements of n pseudocircles is connected for every $n \in \mathbb{N}$.*

Chapter 8

Triangles in Arrangements of Pseudocircles

In Chapter 8.1, we give counterexamples to Grünbaum's conjecture [Grü80, Conjecture 3.7] which are simple. With a recursive construction based on an example with 12 pseudocircles and 16 triangles we obtain a family of digon-free intersecting arrangements with $p_3/n \xrightarrow{n \rightarrow \infty} 16/11 = 1.45$. We then replace Grünbaum's conjecture by Conjecture 8.2: *The lower bound $p_3(\mathcal{A}) \geq 4n/3$ is tight for infinitely many simple arrangements.*

A specific arrangement \mathcal{N}_6^Δ of 6 pseudocircles of 8 triangles is interesting in this context. The arrangement \mathcal{N}_6^Δ has no representation with circles, two different proofs for the non-circularizability of \mathcal{N}_6^Δ were given in Chapter 7. The arrangement \mathcal{N}_6^Δ appears as a subarrangement in all known simple, intersecting, digon-free arrangements with $p_3 < 2n - 4$. This motivates the question, whether indeed Grünbaum's conjecture is true when restricted to intersecting arrangements of circles, see Conjecture 8.1. In Chapter 8.1.1 we discuss arrangements with digons. We give an easy extension of the argument of Snoeyink and Hershberger [SH91] to show that these arrangements contain at least $2n/3$ triangles. All intersecting arrangements known to us have at least $n - 1$ triangles and therefore our Conjecture 8.3 is that $n - 1$ is a tight lower bound for intersecting arrangements with digons.

In Chapter 8.2 we study the maximum number of triangles in arrangements of n pseudocircles. We show an upper bound of order $\frac{4}{3}\binom{n}{2} + O(n)$. For the lower bound construction we glue two arrangements of n pseudolines into an arrangement of n pseudocircles. Since respective arrangements of pseudolines are known, we obtain arrangements of pseudocircles with $\frac{4}{3}\binom{n}{2}$ triangles for $n \equiv 0, 4 \pmod{6}$.

8.1 Intersecting Arrangements with few Triangles

The main result of this section is the following theorem, which disproves Grünbaum's conjecture.

Theorem 8.1. *The minimum number of triangles in digon-free intersecting arrangements of n pseudocircles is*

- (i) 8 for $3 \leq n \leq 6$.
- (ii) $\lceil \frac{4}{3}n \rceil$ for $6 \leq n \leq 14$.
- (iii) $< \frac{16}{11}n$ for all $n = 11k + 1$ with $k \in \mathbb{N}$.

Figures 8.1 and 8.2 show intersecting arrangements with the minimum number of triangles for up to 8 pseudocircles, and Figure 8.3(a) shows an arrangement of 12 pseudocircles with the minimum of 16 triangles. We remark that, in total, there are three non-isomorphic intersecting

arrangements of $n = 8$ pseudocircles with $p_3 = 11$ triangles, these are the smallest counterexamples to Grünbaum's conjecture (cf. Lemma 8.2). We refer to our website [FS] for further examples.




Figure 8.1: Digon-free intersecting arrangements of $n = 3, 4, 5$ circles with $p_3 = 8$ triangles. Triangles (except the outer face) are colored gray.

The basis for Theorem 8.1 was laid by exhaustive computations, which generated all intersecting arrangements of up to $n = 7$ pseudocircles. Starting with the unique intersecting arrangement of two pseudocircles, our program recursively inserted pseudocircles in all possible ways. From the complete enumeration, we know the minimum number of triangles for $n \leq 7$. In the range from 8 to 14, we had to iteratively use arrangements of n pseudocircles with a small number of triangles and digons to generate arrangements of $n + 1$ pseudocircles with the same property. Using this strategy, we found intersecting arrangements with $\lceil 4n/3 \rceil$ triangles for all n in this range. The corresponding lower bound $p_3(\mathcal{A}) \geq 4n/3$ is known from [SH91].




Figure 8.2: Digon-free intersecting arrangements of $n = 6, 7, 8$ pseudocircles with 8, 10, 11 triangles, respectively. The arrangement of $n = 6$ pseudocircles on the left-hand side is named \mathcal{N}_6^Δ .

Another result which we obtained from our computer search is the following: the triangle-minimizing example for $n = 6$ is unique, i.e., there is a unique intersecting arrangement \mathcal{N}_6^Δ of 6 pseudocircles with 8 triangles. In Chapter 7 we gave different proofs for the non-circularizability of \mathcal{N}_6^Δ . Since the arrangement \mathcal{N}_6^Δ appears as a subarrangement of all arrangements with less than $2n - 4$ triangles known to us, the following weakening of Grünbaum's conjecture might be true.

Conjecture 8.1 (Weak Grünbaum Conjecture). *Every digon-free intersecting arrangement of n circles has at least $2n - 4$ triangles.*

If this conjecture was true, it would imply a simple non-circularizability criterion for intersecting arrangements: Any arrangement with $p_3 < 2n - 4$ could directly be classified as non-circularizable.

So far we know that this conjecture is true for all $n \leq 9$. The claim, that we have checked all intersecting arrangements with $p_3(\mathcal{A}) < 2n - 4$ in this range, is justified by the following lemma, which restricts the pairs (p_2, p_3) for which there can exist arrangements of n pseudocircles whose

extensions have $p_3(\mathcal{A}) < 2n - 4$. For example, to get all digon-free intersecting arrangements of $n = 9$ pseudocircles with $p_3 \leq 13$ triangles, we only had to extend intersecting arrangements of $n = 7$ and $n = 8$ pseudocircles with $p_3 + 2p_2 \leq 13$ triangles.

Lemma 8.2. *Let \mathcal{A} be an intersecting arrangement of pseudocircles. Then for every subarrangement \mathcal{A}' of \mathcal{A} we have*

$$p_3(\mathcal{A}') + 2p_2(\mathcal{A}') \leq p_3(\mathcal{A}) + 2p_2(\mathcal{A}).$$

Proof. We show the statement for a subarrangement \mathcal{A}' in which one pseudocircle C is removed from \mathcal{A} . The inequality then follows by iterating the argument. The arrangement \mathcal{A}' partitions the pseudocircle C into arcs. Reinsert these arcs one by one.

Consider a triangle of \mathcal{A}' . After adding an arc, one of the following cases occurs: (1) the triangle remains untouched, or (2) the triangle is split into a triangle and a quadrangle, or (3) a digon is created in the region of the triangle.

Now consider a digon of \mathcal{A}' . After adding an arc, one of the following cases occurs: (1) the digon remains untouched, or (2) there is a new digon inside this digon, or (3) the digon has been split into two triangles. \square

Levi [Lev26] has shown that every arrangement of pseudolines in the real projective plane has at least n triangles. Since arrangements of great-(pseudo)circles are in bijection to arrangements of (pseudo)lines (the bijection is explained in Chapter 8.2.1), it directly follows that every arrangement of great-pseudocircles has at least $2n$ triangles. The next theorem applies the same idea to a superclass of great-pseudocircle arrangements. We think of the theorem as support of the Weak Grünbaum Conjecture (Conjecture 8.1).

Theorem 8.3. *Let \mathcal{A} be an intersecting arrangement of n pseudocircles such that there is a pseudocircle C in \mathcal{A} that separates the two intersection points $C' \cap C''$ of any other two pseudocircles C' and C'' in \mathcal{A} . Then the number of triangles in \mathcal{A} is at least $2n$.*

Proof. Since, for every two pseudocircles C' and C'' distinct from C , the two intersection points of $C' \cap C''$ are separated by the pseudocircle C , the pseudocircle C “partitions” the arrangement \mathcal{A} into two projective arrangements of n pseudolines which lie in the two respective hemispheres. According to Levi [Lev26], there are at least n triangles in each of the two arrangements, thus the original arrangement \mathcal{A} contains at least $2n$ triangles. \square

Felsner and Kriegel [FK99] have shown that every arrangement of n pseudolines in the Euclidean plane has at least $n - 2$ triangles. This can again be turned into a result about triangles in arrangements of pseudocircles.

Theorem 8.4. *Let \mathcal{A} be an intersecting arrangement of n pseudocircles. If \mathcal{A} can be extended by another pseudocircle C such that the pseudocircle C separates the two intersection points $C' \cap C''$ of any other two pseudocircles C' and C'' , then the number of triangles in the original arrangement \mathcal{A} is at least $2n - 4$.*

Proof. Since, for every two pseudocircles C' and C'' distinct from C , the two intersection points of $C' \cap C''$ are separated by C , the pseudocircle C splits the arrangement \mathcal{A} into two Euclidean arrangements of n pseudolines which lie in the two respective hemispheres. According to Felsner and Kriegel [FK99], there are at least $n - 2$ triangles in each of the two arrangements. Since the extending pseudocircle C (which can be considered as the line at infinity in the respective Euclidean pseudoline arrangements) is not incident to any of these triangles, the arrangement \mathcal{A} contains at least $2n - 4$ triangles. \square

We now prepare for the proof of Theorem 8.1(iii), for which we construct a family of (non-circularizable) intersecting arrangements of n pseudocircles with less than $\frac{16}{11}n$ triangles. The basis of the construction is the arrangement \mathcal{A}_{12} of 12 pseudocircles with 16 triangles shown

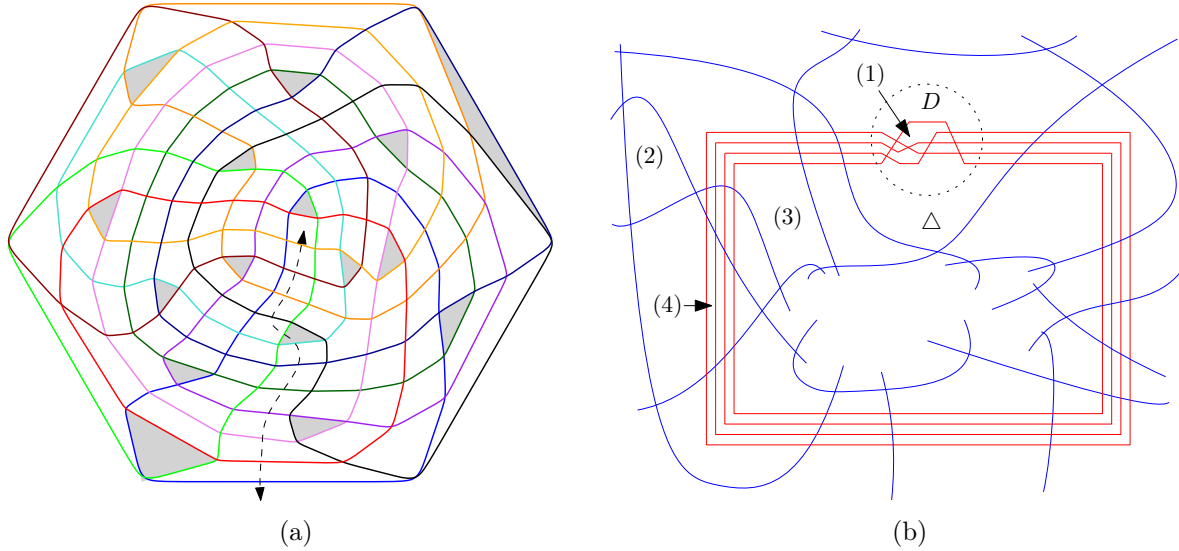


Figure 8.3: (a) The digon-free intersecting arrangement \mathcal{A}_{12} of 12 pseudocircles with exactly 16 triangles. The dashed curve intersects every pseudocircle exactly once. (b) An illustration of the construction in Lemma 8.5. Pseudocircles of \mathcal{A} and \mathcal{B} are drawn red and blue, respectively.

in Figure 8.3(a). This arrangement will be used iteratively for a ‘merge’ as described by the following lemma.

Lemma 8.5. *Let \mathcal{A} and \mathcal{B} be digon-free intersecting arrangements of $n_{\mathcal{A}} \geq 3$ and $n_{\mathcal{B}} \geq 3$ pseudocircles, respectively. If there is a simple curve $P_{\mathcal{A}}$ that*

- (1) *intersects every pseudocircle of \mathcal{A} exactly once*
- (2) *contains no vertex of \mathcal{A} ,*
- (3) *traverses $\tau \geq 1$ triangles of \mathcal{A} , and*
- (4) *forms δ triangles with pairs of pseudocircles from \mathcal{A} ,*

then there is a digon-free intersecting arrangement \mathcal{C} of $n_{\mathcal{A}} + n_{\mathcal{B}} - 1$ pseudocircles with $p_3(\mathcal{C}) = p_3(\mathcal{A}) + p_3(\mathcal{B}) + \delta - \tau - 1$ triangles.

We remark that condition (1) from the statement of Lemma 8.5 asserts that \mathcal{A} is cylindrical. Moreover, if \mathcal{B} is cylindrical, then also \mathcal{C} is cylindrical.

Proof. Take a drawing of \mathcal{A} and make a hole in the two cells which contain the ends of $P_{\mathcal{A}}$. This yields a drawing of \mathcal{A} on a cylinder such that none of the pseudocircles is contractible. The path $P_{\mathcal{A}}$ connects the two boundaries of the cylinder. In fact, the existence of a path with the properties of $P_{\mathcal{A}}$ characterizes cylindrical arrangements.

Stretch the cylindrical drawing such that it becomes a narrow belt, where all intersections of pseudocircles take place in a small disk, which we call *belt-buckle*. This drawing of \mathcal{A} is called a *belt drawing*. The drawing of the red subarrangement in Figure 8.3(b) shows a belt drawing.

Choose a triangle Δ in \mathcal{B} and a pseudocircle B which is incident to Δ . Let b be the edge of B on the boundary of Δ . Specify a disk D , which is traversed by b and disjoint from all other edges of \mathcal{B} . Now replace B by a belt drawing of \mathcal{A} in a small neighborhood of B such that the belt-buckle is drawn within D ; see Figure 8.3(b).

The arrangement \mathcal{C} obtained from *merging* \mathcal{A} and \mathcal{B} , as we just described, has $n_{\mathcal{A}} + n_{\mathcal{B}} - 1$ pseudocircles. Moreover if \mathcal{A} and \mathcal{B} are digon-free/intersecting, then \mathcal{C} has the same property. Most of the cells c of \mathcal{C} are of one of the following four types:

- (1) All boundary edges of c belong to pseudocircles of \mathcal{A} .
- (2) All boundary edges of c belong to pseudocircles of \mathcal{B} .
- (3) All but one of the boundary edges of c belong to pseudocircles of \mathcal{B} and the remaining edge belongs to \mathcal{A} . (These cells correspond to cells of \mathcal{B} with a boundary edge on B .)
- (4) Quadrangular cells, whose boundary edges alternatingly belong to \mathcal{A} and \mathcal{B} .

From the cells of \mathcal{B} , only \triangle and the other cell containing b (which is not a digon since \mathcal{B} is digon-free) have not been taken into account. In \mathcal{C} , the corresponding two cells have at least two boundary edges from \mathcal{B} and at least two from \mathcal{A} . Consequently, neither of the two cells are triangles. The remaining cells of \mathcal{C} are bounded by pseudocircles from \mathcal{A} together with one of the two bounding pseudocircles of \triangle other than B . These two pseudocircles cross through \mathcal{A} following the path prescribed by $P_{\mathcal{A}}$. There are δ triangles among these cells, but τ of these are obtained because $P_{\mathcal{A}}$ traverses a triangle of \mathcal{A} . Among cells of \mathcal{C} of types (1) to (4) all the triangles have a corresponding triangle in \mathcal{A} or \mathcal{B} . But \triangle is a triangle of \mathcal{B} which does not occur in this correspondence. Hence, there are $p_3(\mathcal{A}) + p_3(\mathcal{B}) + \delta - \tau - 1$ triangles in \mathcal{C} . \square

Proof of Theorem 8.1(iii). We use \mathcal{A}_{12} , the arrangement shown in Figure 8.3(a), in the role of \mathcal{A} for our recursive construction. The dashed path in the figure is used as $P_{\mathcal{A}}$ with $\delta = 2$ and $\tau = 1$. Starting with $\mathcal{C}_1 = \mathcal{A}_{12}$ and defining \mathcal{C}_{k+1} as the merge of \mathcal{C}_k and \mathcal{A}_{12} , we construct a sequence $\{\mathcal{C}_k\}_{k \in \mathbb{N}}$ of digon-free intersecting arrangements of $n(\mathcal{C}_k) = 11k + 1$ pseudocircles with $p_3(\mathcal{C}_k) = 16k$ triangles. The fraction $16k/(11k + 1)$ is increasing with k and converges to $16/11 = 1.\overline{45}$ as n goes to ∞ . \square

We remark that using other arrangements from Theorem 8.1(ii) (which also admit a path with $\delta = 2$ and $\tau = 1$) in the recursion, we obtain intersecting arrangements with $p_3 = \lceil \frac{16}{11}n \rceil$ triangles for all $n \geq 6$.

Since the lower bound $\lceil \frac{4}{3}n \rceil$ is tight for $6 \leq n \leq 14$, we believe that the following is true:

Conjecture 8.2. *There exist digon-free intersecting arrangements \mathcal{A} of n pseudocircles with $p_3(\mathcal{A}) = \lceil 4n/3 \rceil$ for infinitely many values of n .*

8.1.1 Intersecting Arrangements with Digons

We know intersecting arrangements of $n \geq 3$ pseudocircles with digons and only $n - 1$ triangles. The examples depicted in Figure 8.4 are part of an infinite family of such arrangements. As illustrated, the intersection order with the black circle determines the arrangement. In fact, it is easy to see that 2^{n-3} different arrangements are possible: Starting with the black, the purple, and the yellow pseudocircles (which give a unique arrangement), each further pseudocircle has its finger placed either immediately to the left or immediately to the right of the previous finger. Figures 8.4(a) and 8.4(b) illustrate the finger-insertion-sequence “right-right-right-...” and “left-right-left-...”, respectively.




Figure 8.4: Intersecting arrangements of n pseudocircles with n digons and $n - 1$ triangles.

Using ideas based on sweeps (cf. [SH91]), we can show that every pseudocircle is incident to at least two triangles. This implies the following theorem:

Theorem 8.6. *Every intersecting arrangement of $n \geq 3$ pseudocircles has at least $2n/3$ triangles.*

The proof of the theorem is based on the following lemma:

Lemma 8.7. *Let C be a pseudocircle in an intersecting arrangement of $n \geq 3$ pseudocircles. Then all digons incident to C lie on the same side of C .*

Proof. Consider a pseudocircle C' that forms a digon D' with C that lies, say, “inside” C . If C'' also forms a digon D'' with C , then C'' has to cross C' in the exterior of C . Hence D'' also has to lie “inside” C . Consequently, all digons incident to C lie on the same side of C . \square

Proof of Theorem 8.6. Let \mathcal{A} be an intersecting arrangement and consider a drawing of \mathcal{A} in the plane. Snoeyink and Hershberger [SH91] have shown that starting with any circle C from \mathcal{A} the outside of C can be swept with a closed curve γ until all of the arrangement is inside of γ . During the sweep γ intersects every pseudocircle from \mathcal{A} at most twice. The sweep uses two types¹ of moves to make progress:

- (1) *take a crossing*, in [SH91] this is called ‘pass a triangle’;
- (2) *leave a pseudocircle*, this is possible when γ and some pseudocircle form a digon which is on the outside of γ , in [SH91] this is called ‘pass a hump’.

Figure 8.5 gives an illustration of the two possible types of moves.




Figure 8.5: An illustration of the two types of moves which are possible in the proof of Theorem 8.6. The blue curve is γ . The interior of γ is left of the shown part of the curve.

Let C be a pseudocircle of \mathcal{A} . By the previous lemma, all digons incident to C lie on the same side of C . Redraw \mathcal{A} so that all digons incident to C are inside C . The first move of a sweep starting at C has to take a crossing, and hence, there is a triangle \triangle incident to C . Redraw \mathcal{A} such that \triangle becomes the unbounded face. Again consider a sweep starting at C . The first move of this sweep reveals a triangle \triangle' incident to C . Since \triangle is not a bounded triangle of the new drawing we have $\triangle \neq \triangle'$, and hence, C is incident to at least two triangles. The proof is completed by double counting the number of incidences of triangles and pseudocircles. \square

Since for $3 \leq n \leq 7$ every intersecting arrangement has at least $n - 1$ triangles, we believe that the following is true:

Conjecture 8.3. *Every intersecting arrangement of $n \geq 3$ pseudocircles has at least $n - 1$ triangles.*

If the arrangement is not required to be intersecting, then the proof of Lemma 8.7 fails. Indeed, if the intersection graph of the arrangement is bipartite, then all faces are of even degree, in particular, there are no triangles; see Figure 8.6(a).

¹There is a third type of move *take a hump* which is the inverse of “leave a pseudocircle”. However, this third type does not occur in the proof of Theorem 8.6 because each two pseudocircles already intersect.




Figure 8.6: Non-intersecting arrangements (a) with no triangles, (b) with a triangle-cell-ratio of $5/6 + O(1/\sqrt{n})$, and (c) with only two non-triangular cells, i.e., with a triangle-cell-ratio of $1 + O(1/n)$.

8.2 Maximum Number of Triangles

Regarding the maximum number of triangles the complete enumeration² provides precise data for $n \leq 8$. Moreover, we used heuristics to generate examples with many triangles for larger n . Table 8.1 summarizes our results and Figures 8.7 and 8.8 show intersecting arrangements of $n = 5, 6, 7, 8$ pseudocircles with the maximal number of triangles; further arrangements are available on our website [FS].

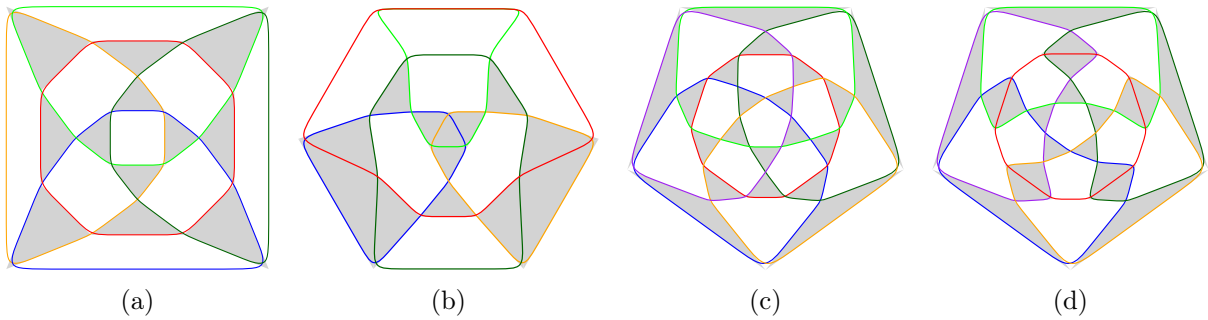


Figure 8.7: (a) and (b) show arrangements of $n = 5$ pseudocircles. The first one is digon-free and has 12 triangles and the second one has 13 triangles and one digon. (c) and (d) show arrangements of $n = 6$ with 20 triangles. The arrangement in (c) is the skeleton of the Icosidodecahedron.

In the next subsection we show that asymptotically the contribution of edges that are incident to two triangles is neglectable. The last subsection gives a construction of intersecting arrangements which show that $\lfloor \frac{4}{3} \binom{n}{2} \rfloor$ is attained for infinitely many values of n .

²While we only have a complete database of arrangements of up to 7 pseudocircles (cf. Table 7.1) we could make sure that there is no arrangement of 8 pseudocircles with at least 39 triangles. Such an arrangement \mathcal{A} would have a pseudocircle C such that the number of triangles of $\mathcal{A}' = \mathcal{A} - C$ would be at least 25. Our computations showed that no arrangement \mathcal{A}' of 7 pseudocircles with $p_3(\mathcal{A}') \geq 25$ can be extended to an arrangement of 8 pseudocircles with more than 38 triangles.




Figure 8.8: (a) An arrangement with $n = 7$ pseudocircles and 29 triangles.
 (b) An arrangement with $n = 8$ pseudocircles and 38 triangles.

n	2	3	4	5	6	7	8	9	10
simple	0	8	8	13	20	29	38	≥ 48	≥ 60
digon-free	-	8	8	12	20	29	38	≥ 48	≥ 60
$\lfloor \frac{4}{3} \binom{n}{2} \rfloor$	1	4	8	13	20	28	37	48	60

Table 8.1: Maximum number of triangles in intersecting arrangements of n pseudocircles.

Recall that we only study simple intersecting arrangements. Grünbaum [Grü80] also looked at non-simple arrangements. His Figures 3.30, 3.31, and 3.32 show drawings of simplicial arrangements that have $n = 7$ with $p_3 = 32$, $n = 8$ with $p_3 = 50$, and $n = 9$ with $p_3 = 62$, respectively. Hence, non-simple arrangements can have more triangles.

Theorem 8.8. *Every intersecting arrangement \mathcal{A} of pseudocircles fulfills $p_3(\mathcal{A}) \leq \frac{4}{3} \binom{n}{2} + O(n)$.*

Proof. Let \mathcal{A} be an intersecting arrangement of $n \geq 4$ pseudocircles. We view \mathcal{A} as a 4-regular plane graph, i.e., the set X of crossings is the vertex set and edges are the segments which connect consecutive crossings on a pseudocircle.

Claim I. *No crossing is incident to 4 triangular cells.*

Assume that a crossing u of C_i and C_j is incident to four triangular cells. Then there is a pseudocircle C_k which bounds those 4 triangles, see Figure 8.9(a). Now C_k only intersects C_i and C_j . This, however, is impossible because $n \geq 4$ and \mathcal{A} is intersecting. \triangle




Figure 8.9: Illustrations of the proof of Claim I and Claim II.

Let $X' \subseteq X$ be the set of crossings of \mathcal{A} that are incident to 3 triangular cells. Our aim is to show that $|X'|$ is small, in fact $|X'| \in O(n)$. When this is shown we can bound the number

of triangles in \mathcal{A} as follows: Under the assumption that $|X'| \in O(n)$, the number of triangles incident to a crossing in X' clearly is in $O(n)$. Now let $Y = X \setminus X'$. Each of the remaining triangles is incident to three elements of Y and each crossing of Y is incident to at most 2 triangles. Hence, there are at most $2|Y|/3 + O(n)$ triangles. Since $|Y| \leq |X| = n(n-1)$ we obtain the bound claimed in the statement of the theorem.

To show that $|X'|$ is small we need some preparation.

Claim II. *Two adjacent crossings u, v in X' share two triangles.*

Since u and v are both incident to 3 triangles, there is at least one triangle Δ incident to both of them. Assume for a contradiction that the other cell which is incident to the segment uv is not a triangle. Let C_i, C_j, C_k be the three pseudocircles such that u is a crossing of C_i and C_j , v is a crossing of C_i and C_k , and Δ is bounded by C_i, C_j, C_k ; see Figure 8.9(b). We denote the third vertex of Δ by w and note that w is a crossing of C_j and C_k .

Since u is incident to three triangles, the segment uw bounds another triangle, which is again defined by C_i, C_j, C_k . Let u' be the third vertex incident to that triangle. Similarly, the segment vw is incident to another triangle which is also defined by C_i, C_j, C_k , and has a third vertex v' .

Again, by the same argument, the segments uu' and vv' , respectively, are both incident to another triangle. However, this is impossible as the two circles C_j and C_k intersect three times. Thus both faces incident to segment uv are triangles. \triangle

Claim III. *Let u, v, w be three distinct crossings in X' . If u is adjacent to both v and w , then v is adjacent to w .*

Since u is incident to three triangles and the segments uv and uw are both incident to two triangles, there is a triangle Δ with corners u, v, w . This triangle shows that u, v , and w are adjacent to each other. \triangle

Claim III implies that each connected component of the graph induced by X' is a complete graph. It is easy to see that a K_4 induced by X' is impossible, and therefore, all components induced by X' are either singletons, edges, or triangles. Figure 8.10 shows the local structure of the arrangement around components of these three types.




Figure 8.10: An illustration of the configurations of crossings in X' . In this figure Δ marks a triangle, “N” marks a k -cell with $k \geq 4$ (“neither a triangle, nor a digon”), “?” marks an arbitrary cell. Crossings with 3 incident triangles are shown as black vertices (these are the crossings in X').

To show that $|X'|$ is small, we are going to trade crossings of X' with digons and then refer to a result of Agarwal et al. [ANP⁺04]. They have shown that the number of digons in intersecting arrangements of pseudocircles is at most linear in n .

To convert crossings of X' into digons we use *triangle flips*. Each of the configurations shown in Figure 8.10 has a gray triangle. By flipping these triangles we obtain the configurations shown in Figure 8.11. These so-obtained configurations have at least as many new digons as the original configurations contain crossings in X' . It may be that the flip creates new triangles and even new vertices which are incident to 3 triangles. However, the flips never remove digons.




Figure 8.11: The configurations in (a), (b), and (c) are obtained by flipping the gray triangle in the configuration from Figure 8.10(a), 8.10(b), and 8.10(c), respectively. The digons created by the flip are marked “D”.

Therefore, thanks to the result from [ANP⁺04] we can make no more than $O(n)$ flips before all the crossings are incident to at most 2 triangles. This finishes the proof of Theorem 8.8. \square

In the proof of Theorem 8.8, we have used flips to trade segments incident to two triangles against digons. It can be shown that at most one component of the graph induced by X' is a K_3 . The proof of this fact is omitted here since it does not improve the bound given in the theorem. Having used a bound on the number of digons we recall that Grünbaum conjectures that $p_2 \leq 2n - 2$ holds for intersecting arrangements.

Since intersecting arrangements have $2\binom{n}{2} + 2$ faces, we can also rewrite the statement of Theorem 8.8: at most $\frac{2}{3} + O(\frac{1}{n})$ of all cells of an intersecting arrangement are triangles. For $n = 7$ there exist arrangements with $29 = \frac{4}{3}\binom{7}{2} + 1$ triangles. It would be interesting to know what the precise maximum value of p_3 for n large is.

For non-intersecting arrangements the arguments from the proof of Theorem 8.8 do not work. Figure 8.6(c) shows an arrangement where all but two cells are triangles. However, if each pseudocircle is required to intersect at least 3 other pseudocircles, then we can proceed similar and show that the triangle-cell-ratio is at most $5/6 + O(1/n)$. In fact, Figure 8.6(b) shows a construction with triangle-cell-ratio $5/6 + O(1/\sqrt{n})$.

Theorem 8.9. *Let \mathcal{A} be an arrangement of n pseudocircles where every pseudocircle intersects at least three other pseudocircles. Then the triangle-cell-ratio is at most $5/6 + O(1/n)$.*

Proof. We proceed as in the proof of Theorem 8.8. In fact, as the “intersecting” property was only used to bound the number of digons, Claims I–III hold also in this less restrictive setting.

From Claims I–III we have learned that every vertex from X' has at least two neighbors from $X \setminus X'$. The following claim will help us to show $|X'| \leq |X \setminus X'|$.

Claim IV. *Every vertex from $X \setminus X'$ has at most two neighbors from X' .*

Suppose for a contradiction that a vertex $v \in X \setminus X'$ has (at least) three neighbors x, y, z from X' . Since x, y, z each have three incident triangular faces and since $v \notin X'$, two of the neighboring faces of v are triangles. In particular, those two triangular faces are not adjacent as otherwise x, y, z would lie in the same component of $G[X']$ and have the same neighbor v – which is impossible.

Without loss of generality, we assume that xy is an edge and that z forms an edge with the fourth neighbor of v , which we denote by w . Since x is incident to a non-triangular face (which is also incident to v), the edge xy bounds another triangle. The same argument shows that zw bounds another triangle, and therefore, the two pseudocircles passing through v intersect three times – a contradiction; see Figure 8.12. This finishes the proof of Claim IV. \triangle




Figure 8.12: An illustration of the proof of Claim IV.

We can now discharge $1/2$ from every vertex of X' to its neighbors from $X \setminus X'$ and, by Claim IV, count at most 1 at each of the vertices from $X \setminus X'$. Therefore, $|X'| \leq |X \setminus X'|$ holds. By counting the face-vertex-incidences we get

$$p_3 \leq \frac{3|X'| + 2|X \setminus X'|}{3} \leq \frac{5|X|}{6},$$

and since the number of faces equals $|X| + 2$, this completes the proof of Theorem 8.9. \square

8.2.1 Constructions using Arrangements of Pseudolines

Great-circles on the sphere are a well known model for projective arrangements of lines (cf. Chapter 6.1). Antipodal pairs of points on the sphere correspond to points of the projective plane. Hence, the great-circle arrangement corresponding to a projective arrangement \mathcal{A} of lines has twice as many vertices, edges, and faces of every type as \mathcal{A} . The same idea can be applied to projective arrangements of pseudolines. If \mathcal{A} is a projective arrangement of pseudolines, take a drawing of \mathcal{A} in the unit disk D such that every line ℓ of \mathcal{A} connects two antipodal points of D . Project D to the upper hemisphere of a sphere S , such that the boundary of D becomes the equator of S . Use a projection through the center of S to copy the drawing from the upper hemisphere to the lower hemisphere of S . By construction the two copies of each pseudoline from \mathcal{A} join together to form a pseudocircle. The collection of these pseudocircles yields an intersecting arrangement of pseudocircles on the sphere with twice as many vertices, edges, and faces of every type as \mathcal{A} . Arrangements of pseudocircles obtained by this construction have a special property:

- If three pseudocircles C , C' , and C'' have no common crossing, then C'' separates the two crossings of C and C' .

Grünbaum [Grü80] calls arrangements with this property ‘symmetric’. In the context of oriented matroids the property is part of the definition of arrangements of pseudocircles [BLW⁺99]. We call arrangements with this property “arrangements of great-pseudocircles” as they generalize the properties of arrangements of great-circles (cf. Chapter 8.2.1).

Arrangements of pseudolines which maximize the number of triangles have been studied intensively. Blanc [Bla11] gives tight upper bounds for the maximum both in the Euclidean and in the projective case and constructs arrangements of pseudolines with $\frac{2}{3}\binom{n}{2} - O(n)$ triangles for every n . In particular, for $n \equiv 0, 4 \pmod{6}$ projective arrangements of straight lines with $\frac{2}{3}\binom{n}{2}$ triangles are known; see also [FG18]. This directly translates to the existence of (1) intersecting arrangements of pseudocircles with $\frac{4}{3}\binom{n}{2} - O(n)$ triangles for every n and (2) intersecting arrangements of circles with $\frac{4}{3}\binom{n}{2}$ triangles for $n \equiv 0, 4 \pmod{6}$. The ‘doubling method’ that has been used for constructions of arrangements of pseudolines with many triangles, see [Bla11], can also be applied for pseudocircles. In fact, in the case of pseudocircles there is more flexibility for applying the method. Therefore, it is conceivable that $\lfloor \frac{4}{3}\binom{n}{2} \rfloor$ triangles can be achieved for all n .

Part III

Discussion

Chapter 9

Discussion

I would like to close this thesis with a discussion on the obtained results, some related results, and with the statement of some open problems.

9.1 Holes in Point Sets

The number of 3-holes and 4-holes ($h_3(n)$ and $h_4(n)$ resp.) is known to be asymptotically quadratic whereas the multiplicative factor of the quadratic term remains unknown. In Chapter 3 we have seen a superlinear lower bound on the number of 5-holes $h_5(n)$, while the currently best known upper bound remains quadratic.

- Are the numbers of 5-holes and 6-holes ($h_5(n)$ and $h_6(n)$ resp.) asymptotically quadratic?
- What are the multiplicative factors of the quadratic terms of $h_3(n)$ and $h_4(n)$?

In particular, we have shown that every ℓ -divided set $S = A \cup B$ with $|A|, |B| \geq 5$ and A, B not in convex position contains at least one ℓ -divided 5-hole (cf. Theorem 3.3). With a generalization of Ham-Sandwich cuts (cf. Lemma 3.4) we then concluded the superlinear bound (cf. Theorem 3.1).

- Is there a way to apply Theorem 3.3 to obtain an $\Omega(n^{1+\epsilon})$ bound on $h_5(n)$?

We remark that all computer-assisted proofs from Chapter 3 towards a proof of Theorem 3.3 were done using C++ and python programs (cf. the flow summary in Figure 3.1). I'm convinced that, when modeling the statement appropriately as a SAT instance (cf. Chapter 4), these results can also be confirmed using SAT solvers. In fact, it might turn out that some of the computer-assisted lemmas can even be shown for larger values of n using SAT solvers, and that some parts of the proof of Theorem 3.3 become obsolete.

Small Point Sets: We have used a SAT model based approach in Chapter 4 to determine the values of $h_5(n)$ for all values $n \leq 16$ (cf. Table 1.1). Due to García ([Gar12], Theorem 3.5) the number of 3-holes and 4-holes only depend on the number of extremal points and on the number of *generated* 3-holes and 4-holes, respectively. Since these values are both bounded from above by $h_5(n)$, the values of $h_3(n)$ and $h_4(n)$ should also be determinable with an appropriate adaption of the SAT model for all values $n \leq 16$.

6-Holes: The minimum number $h(6)$ of points, such that every set of that many points contains a 6-hole, is still unknown today. The best known bounds are $30 \leq h(6) \leq 463$.

At a first glance, it might seem that a similar approach as in Chapter 3 could be used to derive stronger lower bounds also on the minimum number of 6-holes $h_6(n)$. However, since there are point sets of 29 points with no 6-hole [Ove02], one would need to investigate point sets of size at least 30 in order to find an ℓ -divided 6-hole. This task is too demanding for

our implementations, since the number of combinatorially different point sets grows too rapidly. Moreover, the case analysis in several steps of our proof would become much more complicated.

- What is the value of $h(6)$?

We would not be surprised if a SAT model based approach (cf. Chapter 4) can be used to tackle this question.

9.2 Disjoint Holes in Point Sets

In Chapter 4 we used a SAT model based approach to determine the value $h(5, 5) = 17$, that is, every set of 17 points determines two 5-holes with disjoint respective convex hulls. Furthermore, we have investigated interior-disjoint holes.

- It remains open what the values $h(5, 5, 5)$, $h(5, 5, 5, 5)$, \dots are (cf. Chapter 4.2) – and analogously for interior-disjoint holes.

To determine multi-parametric values such as $h(5, 5, 5)$, one can formulate a SAT instance as follows: Three 5-holes X_1, X_2, X_3 are pairwise disjoint if there is a line ℓ_{ij} for every pair X_i, X_j that separates X_i and X_j . By introducing auxiliary variables $Y_{i,j}$ for every pair of 5-tuples X_i, X_j to indicate whether X_i and X_j are disjoint 5-holes, one can formulate an instance in $\Theta(n^{10})$ variables with $\Theta(n^{15})$ constraints. However, since this formulation is quite space consuming, a more compact formulation might be of interest.

Another way to determine the values $h(2, 4, 5)$, $h(3, 4, 5)$, and $h(4, 4, 5)$ is to use partial extension (cf. [Kra03, Chapter 3] and [Sch13]). The idea is to filter order types on 11 points, which do not contain disjoint 2-, 4-, and 5-holes. (To speed up this process, one could first filter order types on 9 points which do not contain disjoint 4- and 5-holes and then extend the remaining order types to 11 points by adding two extremal points.) If no such order type exists, then $h(2, 4, 5)$ would be 11 – otherwise one extends the remaining order types by an additional point outside the convex hull and filters the so-obtained order types on 12 points.

9.3 Erdős–Szekeres type Questions on Colored Point Sets

A k -hole in a point set S is a subset $X \subseteq S$ in convex position with $\text{conv}(X) \cap S = X$ and $|X| = k$. In Chapter 3 we have already seen one generalization of k -holes: A k -island¹ is a subset $X \subseteq S$ (not necessarily in convex position) with $\text{conv}(X) \cap S = X$. An even more general structure is a so-called *generalized k -hole* in a point set S , which is a simple polygon (not necessarily convex) spanned by k points of S with no other point in its interior. Pause to note that every k -island yields at least one generalized k -hole.

Another famous Erdős–Szekeres variant is as follows: Given a set S of n points, each point colored by one of the colors $\{1, \dots, c\}$ for a fixed c , does S contain a monochromatic k -hole, k -island, or generalized k -hole? Devillers and others [DHKS03] showed that the perfect Horton set can be (i) 3-colored such that it does not contain any monochromatic 3-hole, and (ii) 2-colored such that it does not contain any monochromatic 5-hole. Pach and Tóth [PT13] showed that every bicolored point set contains $\Omega(n^{4/3})$ monochromatic 3-holes, and in fact, a quadratic number is conjectured [AFMFP⁺09]. Moreover, it is conjectured that every sufficiently large bicolored point set contains a monochromatic 4-hole [DHKS03]. Some progress was made by Aichholzer et al. [AHH⁺10], who showed that every sufficiently large bicolored set contains at least one monochromatic generalized 4-hole.

¹ k -islands were introduced by Fabila-Monroy and Huemer [FMH12]

- It remains open whether every sufficiently large bicolored set contains monochromatic k -islands for $k \geq 4$ (in particular, monochromatic 4-holes) and monochromatic generalized k -holes for $k \geq 5$.
- Is the number of monochromatic 3-holes in a bicolored set quadratic?

Koshelev [Kos09a] found a set of 46 points without monochromatic 4-holes; see also the point set zoo². Using SAT solvers, we managed to find an abstract order type on 77 points without monochromatic 4-holes. Concerning (realizable) order types, the largest set we found has 48 points.

The largest bicolored set currently known without monochromatic 4-islands has 35 points³. We managed to find an abstract order type on 46 points and a set of 36 points without monochromatic 4-islands.

The largest bicolored set currently known without monochromatic generalized 4-holes has 22 points⁴. We managed to find an abstract order type on 25 points and a set of 24 points without monochromatic generalized 4-holes.

To our knowledge, no bicolored large sets without monochromatic generalized 5-holes were explicitly mentioned in literature (or on auxiliary websites). We found a set of 31 points without monochromatic generalized 5-holes.

All our data is available on our website⁵. We remark that all examples were found using local search techniques such as simulated annealing, SAT solvers, and abstract order type extension (cf. [Kra03, Chapter 2.4]). Mostly all computations were done with our framework pyotlib [Sch14] and, in particular, to find realizations of abstract order types, Wolfram Mathematica [Wol] was used.

9.4 Triangles in Arrangements of Pseudocircles

In Chapter 8 we have seen that Grünbaum's conjecture on triangles in arrangements of pseudocircles turned out wrong in general, however, it might be true when restricted to arrangements of proper circles:

- Is it true that every intersecting digon-free arrangement of n circles has at least $2n - 4$ triangles? (cf. Conjecture 8.1).

Moreover, we have used an intersecting digon-free arrangement of 12 pseudocircles with 16 triangles to recursively construct larger arrangements with $p_3 < \frac{16}{11}n$ triangles.

- We conjecture that arbitrarily large digon-free intersecting arrangements of n pseudocircles with $p_3 = \lceil 4n/3 \rceil$ triangles exist (Conjecture 8.2).

For intersecting arrangements with digons we have shown that they contain at least $2n/3$ triangles. By now we only know of constructions with $n - 1$ triangles, which we verified to be best possible for $n \leq 7$ using computer assistance.

- We conjecture that every intersecting arrangement of n pseudocircles has at least $n - 1$ triangles (cf. Conjecture 8.3).

²http://www.eurogiga-compose.eu/posezo/n46_c2_no_monochromatic_convex_4_hole/n46_c2_no_monochromatic_convex_4_hole.php

³http://www.eurogiga-compose.eu/posezo/n35_c2_no_monochromatic_4_island/n35_c2_no_monochromatic_4_island.php

⁴http://www.eurogiga-compose.eu/posezo/n22_c2_no_monochromatic_4_hole/n22_c2_no_monochromatic_4_hole.php

⁵http://page.math.tu-berlin.de/~scheuch/research/sat_vs_bicolored_point_sets/

9.5 Digons in Arrangements of Pseudocircles

In Chapter 8 we investigated number of triangles in arrangements of pseudocircles. Concerning the number of digons, Grünbaum also has a conjecture:

- Every intersecting arrangement of n pseudocircles contains at most $p_2 \leq 2n - 2$ digons [Grü80, Conjecture 3.6].

Agarwal et al. [ANP⁺04] have shown that the conjecture holds for cylindrical arrangements. Moreover, for intersecting arrangements of pseudocircles they show that the number of digons is at most linear in n , the multiplicative constant of the linear term however remains unknown. Their result is based on the fact, that every intersecting arrangement of n pseudocircles has a constant sized piercing number (later shown to be at most 4 [HPKM⁺18, CKM18]), and hence can be decomposed into constantly many cylindrical subarrangements.

Grünbaum's conjecture is only of interest for intersecting arrangements, since arrangements of n circles exist with a superlinear number of digons⁶. An upper bound of $O(n^{3/2+\epsilon})$ on the number of digons in general arrangements was shown by Aronov and Sharir [AS02].

- What is the maximum number of digons in an arrangement of pseudocircles?

It is worth mentioning that, for the more restrictive setting of arrangements of unit-circles, Pinchasi showed an upper bound of $p_2 \leq n + 3$ [Pin02, Lemma 3.4 and Corollary 3.10].

During our studies we came across the following result⁷:

Proposition 9.1. *Every digon-free arrangement \mathcal{A} of $n \geq 3$ pseudocircles has a Krupp subarrangement.*

Proof. Suppose that \mathcal{A} is digon-free and has no Krupp subarrangement. We obtain a contradiction by constructing an infinite sequence of nested lenses of \mathcal{A} . Let C_0 be an arbitrary pseudocircle of \mathcal{A} . Because there is no Krupp in \mathcal{A} , there is a pseudocircle C_1 such that the two intersections of C_0 and C_1 are consecutive on C_0 along an arc a_0 . The first lense L_1 is the subregion in the interior of C_0 whose boundary consists of a_0 and a piece b_1 of C_1 . Since L_1 is not a digon in \mathcal{A} , there are pseudocircles intersecting the boundary, in fact, they intersect b_1 . Because there is no Krupp involving C_1 , there is a pseudocircle C_2 such that the two intersections of C_1 and C_2 are consecutive on b_1 along an arc a_1 . Lens L_2 is the subregion of L_1 whose boundary consists of a_1 and a piece b_2 of C_2 . Next we find a pseudocircle C_3 whose intersections with C_2 are consecutive on b_2 and so on. \square

It would not be surprising, if the following generalization of the circle-plane representation (cf. Chapter 6) was true:

- Given an arrangement \mathcal{A} of pseudocircles on the sphere \mathbb{S} , does there exist an arrangement \mathcal{E} of pseudoplanes (in \mathbb{R}^3) such that the intersection of \mathcal{E} with \mathbb{S} is precisely the arrangement \mathcal{A} ?

The intersection of two pseudoplanes (which is a pseudoline) would intersect S if and only if the corresponding pseudocircles intersect, and three pseudocircles would form a Krupp if and only if the intersection point of three pseudoplanes lies inside S . Also Proposition 9.1 and many other nice properties would follow directly from the existence such a representation.

⁶As noted by Agarwal and others [ANP⁺04], the bound follows from a construction by Erdős [Erd46] with n circles which has $\Omega(n^{1+c/\log \log n})$ pairs at distance 2.

⁷With a more fine-grained analysis one can show that an arrangement, in which every triple is Non-Krupp, has at least 2 digons, which moreover is best-possible.

9.6 Larger Cells in Arrangements

Besides the number of digons and triangles in arrangements, also larger cells have been investigated (cf. [FG18]). It is known, that infinitely many arrangements without quadrilaterals exist [Rou86], and that arbitrarily large arrangements exist, where all cells are either triangles or quadrangles [LLM⁺07]. Roudneff [Rou87], showed that every arrangement of $n \geq 5$ (pseudo)lines in the projective plane contains linearly many quadrilaterals or linearly many pentagons.

Concerning the upper bound on the number of k -cells in arrangements, Grünbaum [Grü80] showed that every (general) arrangement of n lines determines at most $\frac{1}{2}n(n-3)$ quadrilaterals and conjectured that equality can only hold for simple arrangements. This conjecture was proven to be correct by Roudneff [Rou87], and Forge and Ramírez-Alfonsín [FRA01] showed that the statement is also true for arrangements of pseudolines.

9.7 Circularizability

In the course of Chapter 7, we generated circle representations or proved non-circularizability for all connected arrangements of $n \leq 5$ pseudocircles (cf. Chapter 7.2) and for all digon-free intersecting arrangements of $n \leq 6$ pseudocircles (cf. Chapter 7.2). Besides that, we also investigated the next larger classes and found

- about 4 400 connected digon-free arrangements of 6 circles (which is about 98%),
- about 130 000 intersecting arrangements of 6 circles (which is about 90%), and
- about 2 millions intersecting digon-free arrangements of 7 circles (which is about 66%).

For our computations (especially the last two additional items), we had up to 24 CPUs running over some months with the quantitative realization approaches described in Chapter 7.5.2.

We further investigated arrangements that were not realized by our computer program and have high symmetry or other interesting properties. Non-circularizability proofs for some of these candidates were presented in Chapter 7.4. Since we have no automated procedure for proving non-circularizability, these proofs had to be done by hand.

In case of abstract order types, Bokowski and Richter [BR90] introduced a technique – commonly known as the method of “bi-quadratic final polynomials” – for proving non-realizability, which is based on the Graßmann–Plücker relations (cf. Chapter 2.2). We outline the idea: Assuming that an abstract order type is realized by a point set, then we can find values of the determinants

$$D_{abc} := \det \begin{pmatrix} 1 & 1 & 1 \\ x_a & x_b & x_c \\ y_a & y_b & y_c \end{pmatrix}$$

which clearly have to respect the Graßmann–Plücker relations (cf. Chapter 2.2). In particular, Setting $a_3 = b_3$ in the relation, one of the three summands on the right-hand side of the equation vanishes. After moving negative summands (each is a product of two determinants) to the other side such that all summands are positive, we obtain an equation of the form

$$D_1 \cdot D_2 + D_3 \cdot D_4 = D_5 \cdot D_6,$$

where each of the three products is positive. In particular, since we have $D_1 \cdot D_2 < D_5 \cdot D_6$, we can apply logarithms on both sides and result in the linear inequality $L_1 + L_2 < L_5 + L_6$, where $L_{abc} = \log |D_{abc}|$. In total, we obtain a system of $\Theta(n^5)$ linear inequalities in n^3 variables⁸. It is

⁸To be precise, if the strict inequalities have a solution, then the solution can be scaled such that $L_1 + L_2 + 1 \leq L_5 + L_6$ is fulfilled.

indeed a necessary criterion for realizability that this system is feasible, and checking feasibility of linear programs is known to be doable in polynomial time⁹.

Interestingly, all non-realizable abstract order types on up to $n = 11$ points were certified as non-realizable using this procedure [Kra03]. The smallest non-realizable abstract order type known today which does not violate this necessary criterion has 14 points and is not in general position [RG96].

Since order types correspond to arrangements of great-circles, this non-realizability criterion can also be used to prove non-circularizability for certain arrangements of great-pseudocircles (cf. Corollary 7.6). For general arrangements, however, we do not have such a criterion.

- Is there an efficiently checkable criterion to certify non-circularizability of arrangements of (not necessarily great-)pseudocircles?

It is also highly interesting that Bokowski, Richter, and Sturmfels [BRS90] proposed a more general technique – the method of “final polynomials” – which uses a real version of Hilbert’s Nullstellensatz to certify non-realizability for *all* non-realizable abstract order types. Such final polynomials, however, cannot be computed efficiently.

9.8 Number of Arrangements

As we have seen in Chapter 7.5.4, there are $2^{\Theta(n^2)}$ arrangements on n pseudocircles and only $2^{\Theta(n \log n)}$ arrangements on n circles. This was not surprising as the same numbers apply to arrangements of lines and pseudolines, respectively. Besides the asymptotics, also the constant in the exponent has been studied intensively for arrangements of lines and pseudolines [Fel97, FV11, DM18] (see also [FG18]).

- What are the asymptotics for the various classes of arrangements (connected, intersecting, great-circles, digon-free, cylindrical, ...)?

It is not hard to see that the constant for intersecting arrangements of n pseudocircles is at least twice as large as the constant for arrangements of n great-pseudocircles. For example, take any two arrangements of n pseudolines and glue them to one intersecting arrangement of n pseudocircles.

When investigating the multiplicative constant in the $2^{O(n^2)}$ bound of the number $A(n)$ of (not necessarily connected) arrangements on n pseudocircles, the following approach allows better estimates than the technique from Proposition 7.22: Assuming that the number $C(n)$ of connected arrangements of n pseudocircles is bounded by $C(n) \leq c^n$ for some constant c , then we have

$$A(n) \leq \sum_{\substack{k_1+\dots+k_l=n \\ l, k_1, \dots, k_l \in \mathbb{N}}} \prod_i C(k_i) \leq \sum_{\substack{k_1+\dots+k_l=n \\ l, k_1, \dots, k_l \in \mathbb{N}}} c^{k_1^2+\dots+k_l^2} \leq \sum_{\substack{k_1+\dots+k_l=n \\ l, k_1, \dots, k_l \in \mathbb{N}}} c^{n^2} \leq 2^n \cdot c^{n^2} \leq 2^{O(n^2)}.$$

An analogous argument also applies to arrangements of circles.

9.9 Convexibility

Besides circularizability, Kang and Müller [KM14] have also investigated *convexibility* – the problem of deciding whether an arrangement of pseudocircles is isomorphic to an arrangement

⁹In one of my Bachelor’s theses [Sch14] I’ve developed a python framework for investigating non-realizability of abstract order types using computer assistance. The largest configuration certified non-realizable had 96 points and contained a non-realizable subconfiguration on 17 points, which was found heuristically.

of convex curves¹⁰. They proved that deciding circularizability for connected arrangements is NP-hard¹¹ in general and asked the following question:

- What is the smallest non-convexible arrangement of pseudocircles?

We have convinced ourselves that all (non-circularizable) arrangements of 5 pseudocircles are convexible. Note that stereographic projections do not necessarily map convex curves to convex curves, thus convexibility – in contrast to circularizability – might depend on the planar embedding of the arrangement.

When applying the reduction from [KM14] to the simple non-Pappos arrangement of 9 pseudolines, one obtains a non-convexible arrangement of 18 pseudocircles.

Since ellipses are between circles and convex curves, we also took a short look at “ellipsability”. In fact, we found realizations of all arrangements of $n \leq 5$ pseudocircles with ellipses, which brings us to the following strengthened question:

- What is the smallest non-ellipsable arrangement of pseudocircles?¹²

Moreover,

- is there a generalization of the circle-plane representations to ellipses (involving conic shapes)?

Note that this is a weakening of the question, whether all arrangements can be represented by pseudoplanes (cf. Chapter 9.5)

Another interesting result, which involves cylindrical arrangements and convexibility is the following: Bultena, Grünbaum, and Ruskey [BGR99] showed that every arrangement of closed curves in the plane, with the property that all the curves have a common point in their interior, and that any two curves intersect in finitely many points, admits a convex representation. Since, in planar arrangements of great-pseudocircles, all pseudocircles have a common interior point (the antipodal cell of the outer cell), there exists a convex representation.

Corollary 9.2. *Every planar arrangement of great-pseudocircles is convexible.*

Since great-pseudocircle arrangements are always convexible, and connected arrangements might not be convexible, we pose the following question:

- What classes of arrangements are convexible and for which classes is the decision problem ETR-hard? In particular, are there intersecting arrangement that are non-convexible?

9.10 Extending Arrangements

Besides the Sweeping Theorem for arrangements of pseudocircles, Snoeyink and Hershberger [SH91] also proved the Extension Theorem for arrangements of pseudocircles. It asserts that, given an arrangement \mathcal{A} of pseudocircles and 3 points (not on a common pseudocircle from \mathcal{A}), then there is a pseudocircle C through the 3 points such that $\mathcal{A}+C$ is a pseudocircle arrangement. This theorem can be considered as an analoga of Levi’s Extension Lemma [Lev26], which asserts that, given an arrangement \mathcal{A} of pseudolines and 2 points (not all on a common pseudoline from \mathcal{A}), then there is a pseudoline L through the 2 points such that $\mathcal{A} + L$ is a pseudoline arrangement. It is also worth mentioning that, for $k \geq 3$, Snoeyink and Hershberger found (projective) arrangements of curves that intersect pairwise in at most k points which can neither be swept nor extended.

¹⁰the convexibility problem was already mentioned by Grünbaum [Grü80, page 68]

¹¹actually, they even showed $\exists\mathbb{R}$ -hardness as they give a reduction from stretchability; Figure 36 from [KM14] nicely illustrates the idea

¹²We believed for some time that the Edelsbrunner–Ramos arrangement of 6 pseudocircles is not ellipsable – this turned out to be wrong as we found a realization.

9.11 Cylindrical Arrangements

Recall that cylindrical arrangements have two cells which are separated by each of the pseudocircles of the arrangement. Let \mathcal{A} be a cylindrical arrangement and let z and z' be points in two separated cells, i.e., separated by each of the pseudocircles. The Extension Theorem [SH91] guarantees the existence of a pseudocircle C containing z and z' . In particular, we obtain two disjoint paths from z to z' . From the choice of the two points it follows that each of the two paths has to cross each pseudocircle of \mathcal{A} , hence, each of the two paths crosses each pseudocircle of \mathcal{A} exactly once. Let P be one of the two paths oriented from z to z' . Starting from P , we can sweep the full arrangement and end in P again. If \mathcal{A} has k crossings, the sweep can be formalized as a sequence P_0, \dots, P_k of internally disjoint paths from z to z' with $P_0 = P_k = P$ such that for each crossing c there is an i such that $P_i \cup P_{i+1}$ separates c from all the other crossings.

The order of the crossings along P_i yields a permutation π_i of the n pseudocircles of \mathcal{A} . The sequence π_0, \dots, π_{k-1} can be used to draw a wiring diagram of \mathcal{A} where the wires are n non-intersecting belts on a cylinder. Here a belt is a circle separating the two boundary components of the cylinder.

In particular, we obtain:

Proposition 9.3. *Every cylindrical arrangement \mathcal{A} of pseudocircles has a monotone representation on the cylinder $S^1 \times I$, i.e., for each $x \in S^1$ the fibre I_x has a unique point of intersection with each pseudocircle of \mathcal{A} .*

Similar ideas have been used by Bultena, Grünbaum, and Ruskey [BGR99] to show that every cylindrical arrangement of pseudocircles admits a convex representation, i.e., a representation in the plane where the interior of each pseudocircle is convex.

Another result that we found is the following.

Proposition 9.4. *A connected arrangement on n pseudocircles is cylindrical if and only if its dual graph has diameter n (which is maximal).*

Proof. Given an arrangement \mathcal{A} , choose two points p, p' from cells of maximum distance and apply the Extension Theorem to obtain a pseudocircle C containing p and p' . Since C intersects every pseudocircle of \mathcal{A} at most twice, one of the two arcs of C between p and p' has at most n intersections. Hence, the diameter of the dual graph is at most n .

If \mathcal{A} is cylindrical, then the diameter of the dual graph of \mathcal{A} is clearly n .

On the other hand, if \mathcal{A} is not cylindrical, there is a pseudocircle C in \mathcal{A} which is not separating the two points p and p' . Starting from C sweep the side containing the two points. Eventually the sweep hits the first of the two points, say p . Snoeyink and Hershberger have shown that the sweep always has two options for making progress, therefore, the sweep can be continued with the additional property that p stays on the sweep-front. When the sweep also hits p' , we have a pseudocircle C' containing p and p' and avoiding C . One of the two arcs of C' between p and p' has at most $n - 1$ intersections. Hence, the diameter of \mathcal{A} is less than n . \square

This proposition allows to test efficiently, whether a given arrangement is cylindrical or not: we can simply compute the diameter of the dual graph.

9.12 Flipgraph of Arrangements

In Chapter 7.5.5 we have discussed flip-graphs of various classes of arrangements. In particular, by a continuous motion argument we have shown that the triangle flip-graph of digon-free arrangements of circles is connected. The question whether the same is true for digon-free arrangements of pseudocircles remains open.

- We conjecture that the triangle flip-graph on the set of all intersecting digon-free arrangements of n pseudocircles is connected for every $n \in \mathbb{N}$ (Conjecture 7.2).

9.13 Colorability of Arrangements

It is well known that planar graphs are 4-colorable, hence the arrangement graph of any arrangement of pseudocircles is 4-colorable. Koester [Koe90] constructed a connected arrangement of 7 circles which indeed requires 4 colors, however, it remains unknown whether 3 colors are sufficient to color intersecting arrangements of (pseudo)circles and, in particular, arrangements of great-(pseudo)circles. Hence, we restate a conjecture by Felsner et al. [FHNS06] (see also [FHNS]), which has been verified for $n \leq 11$ [Kra03]:

Conjecture 9.1 ([FHNS06]). *Every arrangement graph of a set of great-circles is 3-colorable.*

9.14 A Generalization of the Erdős–Szekeres Theorem

In Chapter 1 we have seen the classical Erdős–Szekeres Theorem for point sets. Quite recently Medina, Ramírez-Alfonsín, and Salazar [MRAS18] generalized Erdős–Szekeres Theorem to arrangements of intersecting pseudocircles: They use Ramsey-theory to show that every sufficiently large intersecting arrangement contains one of three “unavoidable” arrangements. One of these three types is the famous “cyclic” arrangement (of great-circles), which corresponds to a point set in convex position. Their bound, however, is of order $2^{2^{c \cdot n^2}}$, which is quite far from the (exponential) bound from the original theorem, which leaves the following question open:

- What is the largest relative size of an unavoidable arrangement?

In particular, we wonder whether the piercing constant (cf. [CKM18]) allows an improvement?

9.15 Encoding of Arrangements

Encoding the abstract order type requires quadratically many bits as the number of abstract order types (or pseudoline arrangements) is of order $2^{\Theta(n^2)}$. They can for example be encoded as proposed by Felsner [Fel97], or their planar dual graphs can simply be encoded with quadratically many bits as proposed by Bonichon, Gavaille, and Hanusse [BGH03]. The same clearly also applies to arrangements of pseudocircles.

Concerning the encoding of *realizable* order types, Goodman Pollack, and Sturmfels [GPS89] showed that encoding an order type with integer coordinates might require coordinates of doubly exponential size. This clearly also applies to arrangements of great-circles if they are encoded by planes that intersect the sphere.

- Is there a doubly-exponential lower bound if circles are encoded by their centers and radii (as integer numbers)?

Recently Cardinal et al. [CCI⁺18] presented the first efficient encoding of order types which only uses a subquadratic number of bits. Pause to note that there is a trivial encoding of order types with $\Theta(n \log n)$ bits, for example, by encoding the index of the arrangement in the list of all arrangements. This encoding however is *not* suitable for practical purposes. It is not (known to be) efficiently computable and also triple orientations cannot be decoded efficiently.

- Is there a practical subquadratic encoding for arrangements of circles?

It is also worth mentioning that there exist *degenerate* order types which *cannot* be represented by points with rational coordinates [Grü03, Chapter 5]. We do not know whether this carries over to representations of arrangements, where every circle $C : x^2 + y^2 = r^2$ is described by rational values x, y, r .

9.16 Testing Isomorphism of Arrangements

In general, computing the canonical labeling of a graph is at least as hard as graph isomorphism, however, linear time algorithms are known for planar graphs. Consequently the canonical label for the primal-dual graph and for the dual graph can be computed in $O(n^2)$ time. For our computations, we have used the canonical label function in SageMath [S⁺17a] (see also [S⁺17c]).

9.17 SAT-Encoding of Arrangements

As we have seen in Chapter 4, (abstract) order types can be encoded for a SAT model where variables encode the information of triple orientation and clauses ensure that the correctness (not all \pm assignments to triples correspond to an abstract order type). More precisely, we have used the $\Theta(n^4)$ signotope axioms, which are equivalent to the exchange axioms when assuming that the points are sorted from left to right.

Ortner [Ort08] showed that two arrangements of n pseudocircles are isomorphic if and only if there is suitable permutation of the pseudocircles such that all induced subarrangements of 4 pseudocircles coincide. In fact, he presents two arrangements of 4 pseudocircles which cannot be distinguished only by their induced triples. Since there are only constantly many arrangements of 4 pseudocircles, one can clearly encode arrangements using SAT models with $\Theta(n^4)$ constraints. However, since there are many arrangements on 4 pseudocircles, the native encoding comes with a huge multiplicative constant for the n^4 -term. Hence, I'm not sure whether this model can be used in practice.

9.18 A Generalization of the Zone Theorem

For arrangements of pseudolines (and hence also arrangements of great-circles) the classical Zone theorem (cf. [BEPY91, Pin11]) shows that the complexity of the cells along each pseudoline is of order $O(n)$. Edelsbrunner and others [EGP⁺92] provided a natural generalization of the Zone Theorem: for arrangements of curves with at most k pairwise intersections for some fixed k , they bound the complexity of the cells along each curve by $n \cdot 2^{\text{poly}(\alpha(n))}$ – a function which is almost linear in n . Here $\alpha(n)$ denotes the functional inverse of Ackermann's function. In particular, for an arrangement of n pseudocircles they show that the complexity is of order $O(n \cdot 2^{\alpha(n)})$ and mention that a lower bound of order $\Omega(n \cdot \alpha(n))$ was shown in an unpublished manuscript by Peter Shor¹³

- In which classes of arrangements do the cells along each pseudocircle have a linear complexity?

Moreover, given an arrangement \mathcal{A} of n pseudolines, the Zone Theorem can be used to show $\sum |c|^2 = O(n^2)$, where the sum is over all cells c of \mathcal{A} . Since the generalization of the Zone Theorem from [EGP⁺92] does not have a linear upper bound, the proving technique does not directly apply to arrangements of pseudocircles, which brings us to the following question:

- To which classes of arrangements does the bound $\sum |c|^2 = O(n^2)$ apply?

¹³Such a construction can be obtained as follows: Take an arrangement of n straight-line segments above the x -axis with a lower envelope of superlinear complexity, where moreover all end points of the segments lie in the lower envelope (see e.g. Chapter 7.2 of Matoušek's book [Mat02]). Now take a vertically mirrored copy of all these segments below the x -axis and close the corresponding segments to quadrilaterals, which pairwise intersect at most twice. With an appropriate pseudocircle along the x -axis, one obtains a zone of superlinear complexity.

9.19 Visualization of Arrangements

A large fraction of the figures of arrangements of pseudocircles in this thesis was generated automatically. The programs are written in the mathematical software SageMath [S⁺17a], they are available on demand.

In Chapter 8, where we have investigated triangular cells in arrangements, we visualized intersecting arrangements of pseudocircles as follows: we drew the primal graph using straight-line segments, pseudocircles are colored by distinct colors, and triangles (except the outer face) were filled gray.

Iterated Tutte Embeddings To generate nice aesthetic drawings automatically, we iteratively use weighted Tutte embeddings. We fix a non-digon cell as the outer cell and arrange the vertices of the outer cell as the corners of a regular polygon. Starting with edge-weights all equal to 1, we obtain an ordinary plane Tutte embedding.

For iteration j , we set the weights (force of attraction) of an edge $e = \{u, v\}$ proportional to $p(A(f_1)) + p(A(f_2)) + q(\|u - v\|/j)$ where f_1, f_2 are the faces incident to e , $A(\cdot)$ is the area function, $\|\cdot\|$ is the Euclidean norm, and p, q are suitable monotonically increasing functions from \mathbb{R}^+ to \mathbb{R}^+ (we use $p(x) = x^4$ and $q(x) = x^2/10$).

Intuitively, if the area of a face becomes too large, the weights of its incident edges are increased and will rather be shorter so that the area of the face will also get smaller in the next iteration. It turned out that in some cases the areas of the faces became well balanced but some edges were very short and others long. Therefore we added the dependence on the edge length which is strong at the beginning and decreases with the iterations. The particular choice of the functions was the result of interactive tuning. The iteration is terminated when the change of the weights becomes small enough.

Visualization using Curves On the basis of the straight-line embedding obtained with the Tutte iteration we use splines to smoothen the curves. The details are as follows. First we take a 2-subdivision of the graph, where all subdivision-vertices adjacent to a given vertex v are placed at the same distance $d(v)$ from v . We choose $d(v)$ so that it is at most $1/3$ of the length of an edge incident to v and then use B-splines to visualize the curves. Even though one can draw Bézier curves directly with Sage, we mostly generated ipe files (xml-format, cf. [Che]) so that we can further process the arrangements. Figures 9.1(a) and 9.1(b) show the straight-line and curved drawing of an arrangement of pseudocircles, respectively.




Figure 9.1: (a) Straight-line and (b) curved drawings of the arrangement of great-pseudocircles, which consists of two copies of (c) the non-Pappos arrangement of pseudolines.

Visualization of Arrangements of Pseudolines We also adapted the code to visualize arrangements of pseudolines nicely. One of the lines is considered as the “line at infinity” which is then drawn as a regular polygon. Figure 9.1(c) gives an illustration.

Visualization of Connected Arrangements Since many of the arrangements from Chapter 7 contain digons – which require some additional efforts in the above-described visualization – we decided to visualize the primal-dual graph of those arrangements instead.

Even though the visualization of the primal-dual graph looks somewhat more natural, all k -cells in the arrangement are visualized as polygons of size $2k$ – which is somewhat misleading. As an example consider the rightmost triangle bounded by the green, the orange, and the black pseudocircle in Figure 9.2(b) which actually looks like a quadrangle. Consequently, for Chapter 8, we decided to stick to the former described visualizations of the primal graph.




Figure 9.2: Two drawings of \mathcal{N}_6^Δ : (a) curved primal graph. (b) curved primal-dual graph.

Bibliography

- [AAK02] O. Aichholzer, F. Aurenhammer, and H. Krasser. Enumerating order types for small point sets with applications. *Order*, 19(3):265–281, 2002.
- [ABH⁺17a] O. Aichholzer, M. Balko, T. Hackl, J. Kynčl, I. Parada, M. Scheucher, P. Valtr, and B. Vogtenhuber. A superlinear lower bound on the number of 5-holes. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG’17)*, pages 69–72, 2017.
- [ABH⁺17b] O. Aichholzer, M. Balko, T. Hackl, J. Kynčl, I. Parada, M. Scheucher, P. Valtr, and B. Vogtenhuber. A superlinear lower bound on the number of 5-holes. In B. Aronov and M. J. Katz, editors, *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International Proceedings in Informatics*, pages 8:1–8:16, 2017.
- [ABH⁺19a] O. Aichholzer, M. Balko, T. Hackl, J. Kynčl, I. Parada, M. Scheucher, P. Valtr, and B. Vogtenhuber. A superlinear lower bound on the number of 5-holes. To appear in *Journal of Combinatorial Theory, Series A*, 2019+. arXiv:1703.05253.
- [ABH⁺19b] O. Aichholzer, M. Balko, M. Hoffmann, J. Kynčl, W. Mulzer, I. Parada, A. Pilz, M. Scheucher, P. Valtr, B. Vogtenhuber, and E. Welzl. Minimal representations of order types by geometric graphs. To appear in *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization*, 2019+. arXiv:1908.05124.
- [AFMFP⁺09] O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, and J. Urrutia. Empty monochromatic triangles. *Computational Geometry: Theory and Applications*, 42(9):934–938, 2009.
- [AFMH⁺14] O. Aichholzer, R. Fabila-Monroy, T. Hackl, C. Huemer, A. Pilz, and B. Vogtenhuber. Lower bounds for the number of small convex k -holes. *Computational Geometry: Theory and Applications*, 47(5):605–613, 2014.
- [AGM99] C. Ahrens, G. Gordon, and W. E. McMahon. Convexity and the beta invariant. *Discrete & Computational Geometry*, 22(3):411–424, 1999.
- [AHH⁺10] O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, and B. Vogtenhuber. Large bichromatic point sets admit empty monochromatic 4-gons. *SIAM Journal on Discrete Mathematics*, 23(4):2147–2155, 2010.
- [AHV12] O. Aichholzer, T. Hackl, and B. Vogtenhuber. On 5-gons and 5-holes. In Marquez, Ramos, and Urrutia, editors, *Computational Geometry: XIV Spanish Meeting on Computational Geometry, EGC 2011*, volume 7579 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2012.

- [Aic] O. Aichholzer. Enumerating order types for small point sets with applications. <http://www.ist.tugraz.at/aichholzer/research/rp/triangulations/orderypes/>.
- [Aic09] O. Aichholzer. [Empty] [colored] k -gons. Recent results on some Erdős–Szekeres type problems. In *Proceedings of XIII Encuentros de Geometría Computacional*, pages 43–52, 2009.
- [AK06] O. Aichholzer and H. Krasser. Abstract order type extension and new results on the rectilinear crossing number. *Computational Geometry: Theory and Applications*, 36(1):2–15, 2006.
- [AK16] M. Albenque and K. Knauer. Convexity in partial cubes: The hull number. *Discrete Mathematics*, 339(2):866–876, 2016.
- [AMRS18] A. Arroyo, D. McQuillan, R. B. Richter, and G. Salazar. Levi’s lemma, pseudolinear drawings of K_n , and empty triangles. *Journal of Graph Theory*, 87(4):443–459, 2018.
- [ANP⁺04] P. K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *Journal of the ACM*, 51(2):139–186, 2004.
- [AS02] B. Aronov and M. Sharir. Cutting circles into pseudo-segments and improved bounds for incidences. *Discrete & Computational Geometry*, 28(4):475–490, 2002.
- [AS09] G. Audemard and L. Simon. Predicting learnt clauses quality in modern SAT solvers. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 399–404, 2009.
- [Bal] M. Balko. <http://kam.mff.cuni.cz/~balko/superlinear5Holes>.
- [BD11] B. B. Bhattacharya and S. Das. On the minimum size of a point set containing a 5-hole and a disjoint 4-hole. *Studia Scientiarum Mathematicarum Hungarica*, 48(4):445–457, 2011.
- [BD13] B. B. Bhattacharya and S. Das. Disjoint empty convex pentagons in planar point sets. *Periodica Mathematica Hungarica*, 66(1):73–86, 2013.
- [BEPY91] M. W. Bern, D. Eppstein, P. E. Plassmann, and F. F. Yao. Horizon theorems for lines and polygons. In J. E. Goodman, R. Pollack, and W. L. Steiger, editors, *Discrete & Computational Geometry: Papers from the DIMACS Special Year*, volume 6 of *Series in Discrete Mathematics and Theoretical Computer Science*, pages 45–66. American Mathematical Society, 1991.
- [BF87] I. Bárány and Z. Füredi. Empty simplices in Euclidean space. *Canadian Mathematical Bulletin*, 30(4):436–445, 1987.
- [BFK15] M. Balko, R. Fulek, and J. Kynčl. Crossing numbers and combinatorial characterization of monotone drawings of K_n . *Discrete & Computational Geometry*, 53(1):107–143, 2015.
- [BGH03] N. Bonichon, C. Gavoille, and N. Hanusse. An information-theoretic upper bound of planar graphs using triangulation. In H. Alt and M. Habib, editors, *Annual Symposium on Theoretical Aspects of Computer Science (STACS 2003)*, pages 499–510. Springer, 2003.

- [BGR99] B. Bultena, B. Grünbaum, and F. Ruskey. Convex drawings of intersecting families of simple closed curves. In *Proceedings of the 11th Canadian Conference on Computational Geometry (CCCG'99)*, pages 18–21, 1999.
- [Bie08] A. Biere. PicoSAT essentials. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 4:75–97, 2008.
- [BK01] I. Bárány and Gy. Károlyi. Problems and results around the Erdős–Szekeres convex polygon theorem. In Akiyama, Kano, and Urabe, editors, *Proceedings of the Japanese Conference on Discrete and Computational Geometry (JCDCG 2000)*, volume 2098 of *Lecture Notes in Computer Science*, pages 91–105. Springer, 2001.
- [Bla11] J. Blanc. The best polynomial bounds for the number of triangles in a simple arrangement of n pseudo-lines. In *Geombinatorics*, volume 21, pages 5–17, 2011.
- [BLW⁺99] A. Björner, M. Las Vergnas, N. White, B. Sturmfels, and G. M. Ziegler. *Oriented Matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2 edition, 1999.
- [BMP05] P. Brass, W. Moser, and J. Pach. *Research Problems in Discrete Geometry*. Springer, 2005.
- [BMS17] A. Biniarz, A. Maheshwari, and M. H. M. Smid. Compatible 4-holes in point sets, 2017. arXiv:1706.08105.
- [Bog] A. Bogomolny. Cut-the-knot: Four touching circles.
<http://www.cut-the-knot.org/Curriculum/Geometry/FourTouchingCircles.shtml#Explanation>.
- [BR90] J. Bokowski and J. Richter. On the finding of final polynomials. *European Journal of Combinatorics*, 11(1):21–34, 1990.
- [BRS90] J. Bokowski, J. Richter, and B. Sturmfels. Nonrealizability proofs in computational geometry. *Discrete & Computational Geometry*, 5(4):333–350, 1990.
- [BS89] J. Bokowski and B. Sturmfels. An infinite family of minor-minimal nonrealizable 3-chirotopes. *Mathematische Zeitschrift*, 200(4):583–589, 1989.
- [BV04] I. Bárány and P. Valtr. Planar point sets with a small number of empty convex polygons. *Studia Scientiarum Mathematicarum Hungarica*, 41(2):243–266, 2004.
- [BV17] M. Balko and P. Valtr. A SAT attack on the Erdős–Szekeres conjecture. *European Journal of Combinatorics*, 66:13–23, 2017.
- [Can88] J. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing (STOC '88)*, pages 460–467. Association for Computing Machinery, 1988.
- [CCI⁺18] J. Cardinal, T. M. Chan, J. Iacono, S. Langerman, and A. Ooms. Subquadratic encodings for point configurations. arXiv:1801.01767, 2018.
- [CG98] F. R. K. Chung and R. L. Graham. Forced convex n -gons in the plane. *Discrete & Computational Geometry*, 19(3):367–371, 1998.
- [CGH⁺15] J. Cano, A. García, F. Hurtado, T. Sakai, J. Tejel, and J. Urrutia. Blocking the k -holes of point sets in the plane. *Graphs and Combinatorics*, 31(5):1271–1287, 2015.

- [Che] O. Cheong. The Ipe extensible drawing editor. <http://ipe.otfried.org/>.
- [CKM18] P. Carmi, M. J. Katz, and P. Morin. Stabbing pairwise intersecting disks by four points, 2018. arXiv:1812.06907.
- [Deh87] K. Dehnhardt. *Leere konvexe Vielecke in ebenen Punktmengen*. PhD thesis, TU Braunschweig, Germany, 1987. In German.
- [DHKS03] O. Devillers, F. Hurtado, G. Károlyi, and C. Seara. Chromatic variants of the Erdős–Szekeres theorem on points in convex position. *Computational Geometry*, 26(3):193–208, 2003.
- [DM18] A. Dumitrescu and R. Mandal. New lower bounds for the number of pseudoline arrangements, 2018. arXiv:1809.03619.
- [DTP09] A. Dumitrescu, G. Tóth, and J. Pach. A note on blocking visibility between points. *Geombinatorics*, 19(2):67–73, 2009.
- [Dum00] A. Dumitrescu. Planar sets with few empty convex polygons. *Studia Scientiarum Mathematicarum Hungarica*, 36(1–2):93–109, 2000.
- [Ede87] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer, 1987.
- [EGP⁺92] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane—topology, combinatorics, and algorithms. *Theoretical Computer Science*, 92(2):319–336, 1992.
- [EJ85] P. H. Edelman and R. E. Jamison. The theory of convex geometries. *Geometriae Dedicata*, 19(3):247–270, 1985.
- [ER97] H. Edelsbrunner and E. A. Ramos. Inclusion-exclusion complexes for pseudodisk collections. *Discrete & Computational Geometry*, 17:287–306, 1997.
- [ER00] H. P. Edelman and V. Reiner. Counting the interior points of a point configuration. *Discrete & Computational Geometry*, 23(1):1–13, 2000.
- [Erd46] P. Erdős. On sets of distances of n points. *The American Mathematical Monthly*, 53(5):248–250, 1946.
- [Erd78] P. Erdős. Some more problems on elementary geometry. *Australian Mathematical Society Gazette*, 5:52–54, 1978.
- [ES35] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.
- [Fel97] S. Felsner. On the number of arrangements of pseudolines. *Discrete & Computational Geometry*, 18(3):257–267, 1997.
- [FG18] S. Felsner and J. E. Goodman. Pseudoline Arrangements. In Toth, O’Rourke, and Goodman, editors, *Handbook of Discrete and Computational Geometry*. CRC Press, third edition, 2018.
- [FHNS] S. Felsner, F. Hurtado, M. Noy, and I. Streinu. Webpage: Open problem garden: 3-colourability of arrangements of great circles.
http://www.openproblemgarden.org/op/3_colourability_of_arrangements_of_great_circles.

- [FHNS06] S. Felsner, F. Hurtado, M. Noy, and I. Streinu. Hamiltonicity and colorings of arrangement graphs. *Discrete Applied Mathematics*, 154(17):2470–2483, 2006.
- [FK99] S. Felsner and K. Kriegel. Triangles in Euclidean arrangements. *Discrete & Computational Geometry*, 22(3):429–438, 1999.
- [FMH12] R. Fabila-Monroy and C. Huemer. Covering islands in plane point sets. In A. Márquez, P. Ramos, and J. Urrutia, editors, *Computational Geometry: XIV Spanish Meeting on Computational Geometry, EGC 2011*, volume 7579 of *Lecture Notes in Computer Science*, pages 220–225. Springer, 2012.
- [FPSS12] J. Fox, J. Pach, B. Sudakov, and A. Suk. Erdős–Szekeres-type theorems for monotone paths and convex bodies. *Proceedings of the London Mathematical Society*, 105(5):953–982, 2012.
- [FRA01] D. Forge and J. Ramírez-Alfonsín. On counting the k -face cells of cyclic arrangements. *European Journal of Combinatorics*, 22(3):307–312, 2001.
- [FS] S. Felsner and M. Scheucher. Webpage: Homepage of pseudocircles. <http://www3.math.tu-berlin.de/pseudocircles>.
- [FS17a] S. Felsner and M. Scheucher. Arrangements of pseudocircles: On circularizability. arXiv:1712.02149, 2017.
- [FS17b] S. Felsner and M. Scheucher. Triangles in arrangements of pseudocircles. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG’17)*, pages 225–228, 2017. <http://csconferences.mah.se/eurocg2017/proceedings.pdf>.
- [FS18a] S. Felsner and M. Scheucher. Arrangements of pseudocircles: On circularizability. In *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG’18)*, pages 15:1–15:6, 2018.
- [FS18b] S. Felsner and M. Scheucher. Arrangements of pseudocircles: On circularizability. In T. Biedl and A. Kerren, editors, *Graph Drawing and Network Visualization (GD 2018)*, volume 11282 of *Lecture Notes in Computer Science*, pages 555–568. Springer, 2018.
- [FS18c] S. Felsner and M. Scheucher. Arrangements of pseudocircles: Triangles and drawings. In F. Frati and K.-L. Ma, editors, *Graph Drawing and Network Visualization (GD 2017)*, volume 10692 of *Lecture Notes in Computer Science*, pages 127–139. Springer, 2018.
- [FS19a] S. Felsner and M. Scheucher. Arrangements of pseudocircles: On circularizability. *Discrete & Computational Geometry*, 2019. Ricky Pollack Memorial Issue.
- [FS19b] S. Felsner and M. Scheucher. Arrangements of pseudocircles: Triangles and drawings. To appear in *Discrete & Computational Geometry*, 2019+. arXiv:1708.06449.
- [FV11] S. Felsner and P. Valtr. Coding and counting arrangements of pseudolines. *Discrete & Computational Geometry*, 46(3), 2011.
- [FW00] S. Felsner and H. Weil. A theorem on higher Bruhat orders. *Discrete & Computational Geometry*, 23:121–127, 2000.

- [FW01] S. Felsner and H. Weil. Sweeps, arrangements and signotopes. *Discrete Applied Mathematics*, 109(1):67–94, 2001.
- [Gar12] A. García. A note on the number of empty triangles. In Marquez, Ramos, and Urrutia, editors, *Computational Geometry: XIV Spanish Meeting on Computational Geometry, EGC 2011*, volume 7579 of *Lecture Notes in Computer Science*, pages 249–257. Springer, 2012.
- [Ger08] T. Gerken. Empty convex hexagons in planar point sets. *Discrete & Computational Geometry*, 39(1):239–272, 2008.
- [Goo80] J. E. Goodman. Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Mathematics*, 32(1):27–35, 1980.
- [GP83] J. E. Goodman and R. Pollack. Multidimensional sorting. *SIAM Journal on Computing*, 12(3):484–507, 1983.
- [GP86] J. E. Goodman and R. Pollack. Upper bounds for configurations and polytopes in \mathbb{R}^d . *Discrete & Computational Geometry*, 1(3):219–227, 1986.
- [GP93] J. E. Goodman and R. Pollack. Allowable sequences and order types in discrete and computational geometry. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, pages 103–134. Springer, 1993.
- [GPS89] J. E. Goodman, R. Pollack, and B. Sturmfels. Coordinate representation of order types requires exponential storage. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing (STOC '89)*, pages 405–410. Association for Computing Machinery, 1989.
- [GRS90] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey Theory, Second edition*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., 1990.
- [Grü80] B. Grünbaum. *Arrangements and Spreads*, volume 10 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1972 (reprinted 1980).
- [Grü03] B. Grünbaum. *Convex Polytopes*. Springer, 2003.
- [H⁺19] M. Hohenwarter et al. GeoGebra 6.0, 2019. <http://www.geogebra.org>.
- [Har78] H. Harborth. Konvexe Fünfecke in ebenen Punktmengen. *Elemente der Mathematik*, 33:116–118, 1978. In German.
- [HMPT17] A. F. Holmsen, H. N. Mojarad, J. Pach, and G. Tardos. Two extensions of the Erdős–Szekeres problem, 2017. arXiv:1710.11415.
- [Hof98] P. Hoffman. *The Man Who Loved Only Numbers: The Story of Paul Erdős and the Search for Mathematical Truth*. Hyperion Books, 1998.
- [Hor83] J. Horton. Sets with no empty convex 7-gons. *Canadian Mathematical Bulletin*, 26:482–484, 1983.
- [HPKM⁺18] S. Har-Peled, H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, M. Sharir, and M. Willert. Stabbing Pairwise Intersecting Disks by Five Points. In W.-L. Hsu, D.-T. Lee, and C.-S. Liao, editors, *29th International Symposium on Algorithms and Computation (ISAAC 2018)*, volume 123 of *Leibniz International Proceedings in Informatics*, pages 50:1–50:12, 2018.

- [HU01] K. Hosono and M. Urabe. On the number of disjoint convex quadrilaterals for a planar point set. *Computational Geometry*, 20(3):97–104, 2001.
- [HU05] K. Hosono and M. Urabe. On the minimum size of a point set containing two non-intersecting empty convex polygons. In J. Akiyama, M. Kano, and X. Tan, editors, *Proceedings of the Japanese Conference on Discrete and Computational Geometry (JCDCG 2004)*, volume 3742 of *Lecture Notes in Computer Science*, pages 117–122. Springer, 2005.
- [HU08] K. Hosono and M. Urabe. A minimal planar point set with specified disjoint empty convex subsets. In H. Ito, M. K. Katoh, and Y. Uno, editors, *Kyoto International Conference on Computational Geometry and Graph Theory (KyotoCGGT 2007)*, volume 4535 of *Lecture Notes in Computer Science*, pages 90–100. Springer, 2008.
- [HU18] K. Hosono and M. Urabe. Specified holes with pairwise disjoint interiors in planar point sets. *AKCE International Journal of Graphs and Combinatorics*, 2018. In press.
- [KKS70] J. D. Kalbfleisch, J. G. Kalbfleisch, and R. G. Stanton. A combinatorial problem on convex n -gons. In *Louisiana Conference on Combinatorics, Graph Theory, and Computing*, pages 180–188, 1970.
- [KM88] M. Katchalski and A. Meir. On empty triangles determined by points in the plane. *Acta Mathematica Hungarica*, 51(3):323–328, 1988.
- [KM14] R. J. Kang and T. Müller. Arrangements of pseudocircles and circles. *Discrete & Computational Geometry*, 51:896–925, 2014.
- [Knu92] D. E. Knuth. *Axioms and Hulls*, volume 606 of *Lecture Notes in Computer Science*. Springer, 1992.
- [Koe90] G. Koester. 4-critical 4-valent planar graphs constructed with crowns. *Mathematica Scandinavica*, 67(1):15–22, 1990.
- [Kos09a] V. Koshelev. On Erdős–Szekeres problem and related problems, 2009. arXiv:0910.2700.
- [Kos09b] V. A. Koshelev. On Erdős–Szekeres problem for empty hexagons in the plane. *Modelirovanie i Analiz Informatsionnykh Sistem*, 16(2):22–74, 2009. In Russian.
- [KP98] D. Kleitman and L. Pachter. Finding convex sets among points in the plane. *Discrete & Computational Geometry*, 19(3):405–410, 1998.
- [Kra03] H. Krasser. *Order Types of Point Sets in the Plane*. PhD thesis, Institute for Theoretical Computer Science, Graz University of Technology, Austria, 2003.
- [Lev26] F. Levi. Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse*, 78:256–267, 1926.
- [LLM⁺07] J. Leanos, M. Lomeli, C. Merino, G. Salazar, and J. Urrutia. Simple Euclidean arrangements with no (≥ 5) -gons. *Discrete & Computational Geometry*, 38(3):595–603, 2007.
- [LO05] J. Linhart and R. Ortner. An arrangement of pseudocircles not realizable with circles. *Beiträge zur Algebra und Geometrie*, 46(2):351–356, 2005.

- [Mat02] J. Matoušek. *Lectures on Discrete Geometry*. Springer, 2002.
- [Mat14] J. Matoušek. Intersection graphs of segments and $\exists\mathbb{R}$, 2014. arXiv:1406.2636.
- [Mil64] J. Milnor. On the Betti numbers of real varieties. *Proceedings of the AMS*, 15:275–280, 1964.
- [Mnë88] N. E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In O. Y. Viro and A. M. Vershik, editors, *Topology and Geometry – Rohlin Seminar*, volume 1346 of *Lecture Notes in Mathematics*, pages 527–543. Springer, 1988.
- [MRAS18] C. Medina, J. Ramírez-Alfonsín, and G. Salazar. The unavoidable arrangements of pseudocircles, 2018. arXiv:1805.10957.
- [MS00] W. Morris and V. Soltan. The Erdős–Szekeres problem on points in convex position – a survey. *Bulletin of the AMS, New Series*, 37(4):437–458, 2000.
- [MV16] H. N. Mojarad and G. Vlachos. An improved upper bound for the Erdős–Szekeres conjecture. *Discrete & Computational Geometry*, 56(1):165–180, 2016.
- [Nic07] M. C. Nicolas. The empty hexagon theorem. *Discrete & Computational Geometry*, 38(2):389–397, 2007.
- [NY16] S. Norin and Y. Yuditsky. Erdős–Szekeres without induction. *Discrete & Computational Geometry*, 55(4):963–971, 2016.
- [O’R94] J. O’Rourke. *Computational Geometry in C*. Cambridge University Press, 1994.
- [Ort08] R. Ortner. Embeddability of arrangements of pseudocircles into the sphere. *European Journal of Combinatorics*, 29(2):457–469, 2008.
- [Ove02] M. Overmars. Finding sets of points without empty convex 6-gons. *Discrete & Computational Geometry*, 29(1):153–158, 2002.
- [Pil18] A. Pilz. A note on the flip distance problem for edge-labeled triangulations, 2018. arXiv:1808.03126.
- [Pin02] Pinchasi. Gallai–Sylvester theorem for pairwise intersecting unit circles. *Discrete & Computational Geometry*, 28(4):607–624, 2002.
- [Pin11] R. Pinchasi. The zone theorem revisited, 2011. http://www2.math.technion.ac.il/~room/ps_files/zonespl.pdf.
- [PO49] I. G. Petrovskii and O. A. Oleĭnik. On the topology of real algebraic surfaces. *Izvestiya Akad. Nauk SSSR. Ser. Mat.*, 13:389–402, 1949. In Russian.
- [PRS06] R. Pinchasi, R. Radoičić, and M. Sharir. On empty convex polygons in a planar point set. *Journal of Combinatorial Theory, Series A*, 113(3):385–419, 2006.
- [PT13] J. Pach and G. Tóth. Monochromatic empty triangles in two-colored point sets. *Discrete Applied Mathematics*, 161(9):1259–1261, 2013.
- [RG95] J. Richter-Gebert. Mnëv’s universality theorem revisited. *Séminaire Lotharingien de Combinatoire*, 34:15 pages, 1995.
- [RG96] J. Richter-Gebert. Two interesting oriented matroids. In *Documenta Mathematica*, volume 1, pages 137–148. Deutsche Mathematiker-Vereinigung, 1996.

- [RG11] J. Richter-Gebert. *Perspectives on Projective Geometry – A Guided Tour through Real and Complex Geometry*. Springer, 2011.
- [Rou86] J.-P. Roudneff. On the number of triangles in simple arrangements of pseudolines in the real projective plane. *Discrete Mathematics*, 60:243–251, 1986.
- [Rou87] J.-P. Roudneff. Quadrilaterals and pentagons in arrangements of lines. *Geometriae Dedicata*, 23(2):221–227, 1987.
- [S⁺17a] W. A. Stein et al. *Sage Mathematics Software (Version 8.0)*, 2017. <http://www.sagemath.org>.
- [S⁺17b] W. A. Stein et al. *Sage Reference Manual: Algebraic Numbers and Number Fields (Release 8.0)*, 2017.
http://doc.sagemath.org/pdf/en/reference/number_fields/number_fields.pdf.
- [S⁺17c] W. A. Stein et al. *Sage Reference Manual: Graph Theory (Release 8.0)*, 2017.
<http://doc.sagemath.org/pdf/en/reference/graphs/graphs.pdf>.
- [Scha] M. Scheucher. <http://www.ist.tugraz.at/scheucher/5holes>.
- [Schb] M. Scheucher. Webpage: On disjoint holes in point sets.
http://page.math.tu-berlin.de/~scheuch/supplemental/5holes/disjoint_holes/.
- [Sch03] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [Sch13] M. Scheucher. *Counting Convex 5-Holes*. Bachelor’s thesis, Graz University of Technology, Austria, 2013. In German.
http://page.math.tu-berlin.de/~scheuch/publ/bachelors_thesis_2013.pdf.
- [Sch14] M. Scheucher. *On Order Types, Projective Classes, and Realizations*. Bachelor’s thesis, Graz University of Technology, Austria, 2014.
http://page.math.tu-berlin.de/~scheuch/publ/bachelors_thesis_tm_2014.pdf.
- [Sch18] M. Scheucher. On disjoint holes in point sets. arXiv:1807.10848, 2018.
- [Sch19a] M. Scheucher. On disjoint holes in point sets. In *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG’19)*, pages 22:1–22:8, 2019.
- [Sch19b] M. Scheucher. On disjoint holes in point sets. In *Acta Mathematica Universitatis Comenianae*, volume 88, pages 1049–1059, 2019.
- [SH91] J. Snoeyink and J. Hershberger. Sweeping arrangements of curves. In J. E. Goodman, R. Pollack, and W. L. Steiger, editors, *Discrete & Computational Geometry: Papers from the DIMACS Special Year*, volume 6 of *Series in Discrete Mathematics and Theoretical Computer Science*, pages 309–349. American Mathematical Society, 1991.
- [Slo] N. J. A. Sloane. The on-line encyclopedia of integer sequences.
<http://oeis.org>.

- [SP06] G. Szekeres and L. Peters. Computer solution to the 17-point Erdős–Szekeres problem. *Australia and New Zealand Industrial and Applied Mathematics*, 48(2):151–164, 2006.
- [SŠ17] M. Schaefer and D. Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60(2):172–193, 2017.
- [SSS19a] M. Scheucher, H. Schrezenmaier, and R. Steiner. A note on universal point sets for planar graphs. In *Proceedings of the 35th European Workshop on Computational Geometry (EuroCG’19)*, pages 21:1–21:9, 2019.
- [SSS19b] M. Scheucher, H. Schrezenmaier, and R. Steiner. A note on universal point sets for planar graphs. To appear in *Proceedings of the 27th International Symposium on Graph Drawing and Network Visualization*, 2019+. arXiv:1811.06482.
- [SU07] T. Sakai and J. Urrutia. Covering the convex quadrilaterals of point sets. *Graphs and Combinatorics*, 23(1):343–357, 2007.
- [Suk17] A. Suk. On the Erdős–Szekeres convex polygon problem. *Journal of the AMS*, 30:1047–1053, 2017.
- [Suv88] P. Suvorov. Isotopic but not rigidly isotopic plane systems of straight lines. In *Topology and Geometry – Rohlin Seminar*, volume 1346 of *Lecture Notes in Mathematics*, pages 545–556. Springer, 1988.
- [SZ10] W. Steiger and J. Zhao. Generalized ham-sandwich cuts. *Discrete & Computational Geometry*, 44(3):535–545, 2010.
- [Tho65] R. Thom. *Sur L’Homologie des Variétés Algébriques Réelles*, pages 255–265. Princeton University Press, 1965.
- [Tut62] W. T. Tutte. A census of planar triangulations. *Canadian Journal of Mathematics*, 14:21–38, 1962.
- [TV98] G. Tóth and P. Valtr. Note on the Erdős–Szekeres theorem. *Discrete & Computational Geometry*, 19(3):457–459, 1998.
- [TV05] G. Tóth and P. Valtr. The Erdős–Szekeres theorem: Upper bounds and related results. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*, volume 52, pages 557–568. MSRI Publications, Cambridge Univ. Press, 2005.
- [Val92a] P. Valtr. Convex independent sets and 7-holes in restricted planar point sets. *Discrete & Computational Geometry*, 7(2):135–152, 1992.
- [Val92b] P. Valtr. Sets in \mathbb{R}^d with no large empty convex subsets. *Discrete Mathematics*, 108(1):115–124, 1992.
- [Val95] P. Valtr. On the minimum number of empty polygons in planar point sets. *Studia Scientiarum Mathematicarum Hungarica*, pages 155–163, 1995.
- [Val08] P. Valtr. On empty hexagons. In *Surveys on Discrete and Computational Geometry: Twenty Years Later*, volume 453 of *Contemporary Mathematics*, pages 433–441. American Mathematical Society, 2008.
- [Val12] P. Valtr. On empty pentagons and hexagons in planar point sets. In *Proceedings of Computing: The Eighteenth Australasian Theory Symposium (CATS 2012)*, pages 47–48, Melbourne, Australia, 2012.

- [War68] H. Warren. Lower bounds for approximation by nonlinear manifolds. *Transactions of the AMS*, 133:167–178, 1968.
- [WHH14] N. Wetzler, M. J. H. Heule, and W. A. Hunt. DRAT-trim: Efficient checking and trimming using expressive clausal proofs. In C. Sinz and U. Egly, editors, *Theory and Applications of Satisfiability Testing – SAT 2014*, volume 8561 of *Lecture Notes in Computer Science*, pages 422–429. Springer, 2014.
- [Wika] Wikipedia contributors. Bundle theorem.
http://en.wikipedia.org/wiki/Bundle_theorem.
- [Wikb] Wikipedia contributors. Cross product generalizations.
http://en.wikipedia.org/wiki/Cross_product#Multilinear_algebra.
- [Wikc] Wikipedia contributors. Existential theory of the reals.
http://en.wikipedia.org/wiki/Existential_theory_of_the_reals.
- [Wol] Wolfram Research Inc. Mathematica 11.3.0.
<http://www.wolfram.com/mathematica>.
- [WW19] U. Wagner and E. Welzl. Connectivity of triangulation flip graphs in the plane. To appear in the Proceedings of 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2020), 2019+.
- [YW15] X. S. You and X. L. Wei. On the minimum size of a point set containing a 5-hole and double disjoint 3-holes. *Mathematical Notes*, 97(5):951–960, 2015.