Transversal Helly numbers, pinning theorems and projection of simplicial complexes

Xavier Goaoc

École doctorale IAEM
Département Informatique

TRANSVERSAL HELLY NUMBERS, PINNING THEOREMS AND PROJECTIONS OF SIMPLICIAL COMPLEXES

Habilitation Thesis

XAVIER GOAOC

Defended on Decembre 7th, 2011

Committee

President		
	Daniel BARLET	Professeur, Université Henri Poincaré-Nancy 1
Reading & defense		
	Jean-Daniel BOISSONNAT	Directeur de recherche, INRIA Sophia-Antipolis
	Micha SHARIR	Professeur, Tel-Aviv University
	Frank SOTTILE	Professeur, Texas A & M University
Defense		
	Horatiu CIRSTEA	Professeur, Université Henri Poincaré-Nancy 1
	Michel POCCHIOLA	Professeur, Université Pierre et Marie Curie-Paris 6
	Jorge-Luis RAMIREZ-ALFONSIN	Professeur, Université Montpellier 2
	Raimund SEIDEL	Professeur, Universität des Saarlandes

FOREWORD

In the french academic system, the habilitation thesis is usually the occasion to stop and look back at the work done over some extended period of time, sometimes since the PhD thesis. As this manuscript focuses on contributions in a particular area, it will only give a partial account of this work. This preamble fills this gap by giving a broader overview of my research activities of these last few years.

⋆ ⋆ ⋆

My main area of research is *discrete and computational geometry*. Discrete geometry is an area of mathematics that studies combinatorial properties, such as packing, covering or incidences, of geometric objects (points, lines, balls, polytopes, lattices...). Computational geometry is an area of theoretical computer science that focuses on algorithms for solving geometric problems; the emphasis is usually put on provably correct algorithms and their complexity analysis. The close interaction between the two fields makes it difficult to draw the line between them: algorithmic considerations inspire geometric questions, and conversely.

An important part of my research activity can be traced back to questions in *line geometry*. Line geometry is a classical subject that has been the focus of ongoing research since the 19th century and provides the foundation for the resolution of several algorithmic questions such as three-dimensional visibility, shape approximation, and regression depth computation. These foundations are, however, incomplete and a number of results need to be refined. This is especially true when it comes to analyzing how known results can be improved by taking into account the shape of the geometric objects defining the sets of lines considered. The results presented in this manuscript are mostly from this line of research [1, 7, 8, 11, 12, 13, 14, 15, 26], which lies at the crossroads of computational geometry, discrete geometry, enumerative geometry, real algebraic geometry and non-linear computational geometry.

Image synthesis and analysis often assumes that light travel along straight lines, making computer graphics and computer vision two natural application areas for line geometry. I worked on two questions originating from these areas and involving, at some level, line geometry: *shadow boundary computation* [Bat08, Dem08, 22, Jan10] (we studied a topology-based refinement of the classical notion of visual event, showing that it suffices to determine shadow boundaries and leads to substantially smaller data structures) and *geometric models for imaging systems* [4] (we generalized two existing geometric models for non-central cameras, establishing their equivalence along the way, and extended to several non-central imaging devices techniques developed for the central camera, e.g. simple ray-shooting and stereo-reconstruction).

The geometric aspects of a geometric problem can sometimes be encapsulated in a few key properties so that what remains is essentially a combinatorial question. Natural bridges therefore appeared between line geometry and combinatorics and combinatorial geometry. This led me to explore questions such as *minimal approximate coverings* [20, 21] (we showed that any complete covering of a convex shape by other convex shape of similar size contains small approximate covers, where the meaning of "small" and "approximate" can be quantified, and identified distinct behaviors depending on the smoothness of the covering shapes), *shatter functions of hypergraphs and families of permutations* [16] (exploring how upper bounds on the size of "projections" of a combinatorial structure on small subsets imply systematic asymptotic upper bounds on the size of the structure) and *projections of simplicial complexes and posets* [19] (which are developed in Chapters 8 and 9).

In computational geometry, the worst-case bounds are usually realized by pathological constructions that seldom occur in practice due to structure in the input or finite precision in its representation. A natural question is to provide more adequate bounds via probabilistic geometric models. In this direction, I worked on the *expected size of 3D Delaunay triangulations* [23] (we extended to cylinders a complexity analysis that previously held for other surfaces, requiring the development of a new set of techniques) and started investigating the *smoothed complexity of convex hulls and other geometric structures* [2] (where we explored how quickly the expected complexity of the convex hull of a set of points drops when the points are perturbed; we obtained near-tight estimates for several models of perturbation and related these estimates to empirical observations of numerical rounding phenomena).

I also had opportunities to work in a few other directions such as *bounded curvature path planning* [25] (we reduced to convex optimization a problem for which only constant-factor approximations were previously known) and *untangling questions in graph drawing* [27, 28] (we established the hardness of minimizing the number of vertex-moves required to turn a given, non-plane, straight-line embedding of a planar graph into a plane graph).

⋆ ⋆ ⋆

Most of the results presented in this thesis were previously published in collaboration with various co-authors. Rather than keeping the structure induced by these publications, I reorganized the material and took advantage of a better hindsight to simplify certain proofs along the way. Specifically, the publications or preprints on which this manuscript is based are:

- *Line transversals to disjoint balls* [7, 8], with Ciprian Borcea and Sylvain Petitjean (Chapter 2),
- *Hadwiger and Helly-type theorems for disjoint unit spheres* [11, 13], with Otfried Cheong, Andreas Holmsen and Sylvain Petitjean (Chapters 2, 3 and 4),
- *Lower bounds to Helly numbers of line transversals to disjoint congruent balls* [12] with Otfried Cheong and Andreas Holmsen(Chapter 3),
- *Geometric permutations of disjoint unit spheres* [14, 15], with Otfried Cheong and Hyeon-Suk Na (Chapter 4),
- *Lines pinning lines* [1] with Boris Aronov, Otfried Cheong and Günter Rote (Chapters 5 and 6),
- *Pinning a Line by Balls or Ovaloids in* $\mathbb{R}^3$ [26] with Stefan König and Sylvain Petitjean (Chapter 7),
- *Helly numbers of acyclic families* [19] with Éric Colin de Verdière and Grégory Ginot (Chapter 9).

This manuscript also incorporates insight from on-going works with Otfried Cheong, Jae-Soon Ha and Jungwoo Yang (Chapter 4), Guillaume Batog (Chapter 7) and Éric Colin de Verdière and Grégory Ginot (Chapters 8).

ABSTRACT

The efficient resolution of various problems in computational geometry, for instance visibility computation or shape approximation, raises new questions in line geometry, a classical area going back to the mid-19th century. This thesis fits into this theme, and studies Helly numbers of certain sets of lines, an index related to certain *basis theorems* arising in computational geometry and combinatorial optimization.

Formally, the *Helly number* of a family of sets with empty intersection is the size of its largest inclusion-wise minimal sub-family with empty intersection. For $d \geqslant 2$ let $\mathcal{H}_d$ denote the least integer such that for any family $\{B_1, \ldots, B_n\}$ of pairwise disjoint balls of equal radius in $\mathbb{R}^d$, the Helly number of $\{\mathcal{T}(B_1), \ldots, \mathcal{T}(B_n)\}$ is at most $\mathcal{H}_d$, where $\mathcal{T}(B_i)$ denotes the set of lines intersecting B_i. In 1957, Ludwig Danzer showed that $\mathcal{H}_2$ equals 5 and conjectured that $\mathcal{H}_d$ is finite for all $d \geqslant 2$ and increases with d. We establish that $\mathcal{H}_d$ is at least $2d-1$ and at most $4d-1$ for any $d \geqslant 2$, proving the first conjecture and providing evidence in support of the second one.

To study Danzer's conjectures, we introduce the *pinning number*, a local analogue of the Helly number that is related to *grasping* questions studied in robotics. We further show that pinning numbers can be bounded for sufficiently generic families of polyhedra or ovaloids in $\mathbb{R}^3$, two situations where Helly numbers can be arbitrarily large.

A theorem of Tverberg asserts that when $\{B_1, \ldots, B_n\}$ are disjoint translates of a convex figure in the plane, the Helly number of $\{\mathcal{T}(B_1), \ldots, \mathcal{T}(B_n)\}$ is at most 5. Although quite different, both our and Tverberg's proofs use, in some way, that the intersection of at least two $\mathcal{T}(B_i)$'s has a bounded number of connected components, each contractible. Using considerations on homology of projection of simplicial complexes and posets, we unify the two proofs and show that such topological condition suffice to ensure explicit bounds on Helly numbers.

CONTENTS

1 INTRODUCTION

Isn't it exciting?
Hazel

This habilitation thesis discusses Helly numbers, a notion that originates in a classical theorem of Eduard Helly often referred to as "one of the pillars of convex geometry" (the other two pillars being Radon's and Caratheodory's theorems).

Theorem 1.1 (Helly's theorem [Hel23])**.** *If any* $d+1$ *members of a finite family of convex sets in* $\mathbb{R}^d$ *have a point in common then the whole family has a point in common.*

In the contrapositive, Helly's theorem states that if finitely many convex sets have empty intersection then a small number of them, at most $d+1$, must already have empty intersection. Helly's theorem initiated a search for conditions that, like convexity, ensure that empty intersection can be witnessed by small subfamilies. Helly's name remained attached to the notion, and the *Helly number* of a family $\mathcal{C}$ of sets is defined as the largest integer k such that there exist $x_1, \ldots, x_k$ in $\mathcal{C}$ satisfying two conditions:

(i) $\bigcap_{1 \leqslant i \leqslant k} x_i = \emptyset$, and

(ii) for any $j \in \{1, \ldots, k\}$, $\bigcap_{1 \leqslant i \leqslant k; i \neq j} x_i \neq \emptyset$.

We are interested in Helly numbers of sets of lines. More precisely, we investigate conditions on families $\{A_1, \ldots, A_n\}$ of subsets of $\mathbb{R}^d$ ensuring that $\{\mathcal{T}(A_1), \ldots, \mathcal{T}(A_n)\}$ has bounded Helly number, where $\mathcal{T}(A_i)$ denotes the set of lines intersecting A_i. This leads to exploring how the geometry of $A_1, \ldots, A_k$ influences the structure of the set $\mathcal{T}(A_1) \cap \ldots \cap \mathcal{T}(A_k)$ of their line transversals. The natural setting to study these questions is *line geometry*, the theory of the space of lines as studied in the second half of the 19th century by people such as Plücker, Klein, Grassmann...

⋆ ⋆ ⋆

Helly numbers come in various guise in mathematics and computer science. Before we outline our results we illustrate this diversity with a few examples.

The area where Helly numbers received most attention is perhaps *discrete geometry*. There are, naturally, a great many Helly numbers following from *direct* geometric conditions on the family under consideration; for instance, the Helly number of any finite family of homothets of a planar convex curve is known to be at most 4 [Swa03]. Many Helly numbers are also known for sets *induced* by simple geometric objects. Let us give an elementary example. Start with a family $\mathcal{F} = \{p_1, \dots, p_n\}$ of points in $\mathbb{R}^d$ and let B_i denote the set of balls of radius r containing p_i. Since the set of centers of balls in B_i is simply the ball with center p_i and radius r, it follows from Helly's theorem that $\{B_1, \dots, B_n\}$ has Helly number at most $d+1$; in other words, if $\mathcal{F}$ cannot be enclosed in a ball of radius r then some $d+1$ points of $\mathcal{F}$ already do not fit in such a ball. Examples of similar Helly numbers include the existence of separating surfaces [Lay72], the possibility of illuminating a region using few light sources [Bre92] or the dimension of the kernel of a polygon [Bre81, Bre03]. Perhaps more surprisingly, there also exist Helly numbers relative to the existence of geometric structures satisfying certain conditions, for instance Minkowski structures making a given set equilateral [Pet71]. Helly numbers were extended in various ways (e.g. via *fractional* or *colorful* analogues) and studied in tropical geometry [GS08, GM10] or more abstract settings such as convexity spaces [Kol91] or matroids [Edm01]. We refer to the classical surveys of Danzer et al. [DGK63], Eckhoff [Eck93] and Wenger [Wen04] for a more detailed account on the study of Helly numbers in discrete geometry.

In *algorithms*, Helly numbers naturally arise in the context of optimization problems where the goal is to maximize (or minimize) some function ϕ under a family $\mathcal{F}$ of constraints. In many situations ϕ takes its value over some geometric space and each constraint requires that the solution belongs to some subset of that space; for instance, in *linear programming* ϕ is defined over $\mathbb{R}^d$ and each constraint requires that the solution lies in a given halfspace. In the case where all constraints cannot be simultaneously satisfied, the maximum size of a *certificate of infeasibility* is naturally given by the Helly number of $\mathcal{F}$. This was, for instance, the motivation for studying Helly numbers in hybrid discrete-continuous settings [AW10]. Helly numbers are also relevant when the problem is feasible. For instance, in the *LP-type problems* framework [SW92], which generalizes linear programming and captures problems such as computing the smallest enclosing ball or cylinder, the complexity of computing a solution depends on the so-called *combinatorial dimension* which is, essentially, the Helly number of the level sets of the function to be optimized [Ame94]. In the discrete realm, a similar setting extends *integer linear programming* and enjoys the same connection to Helly numbers [Hal04]. Other applications of Helly numbers in optimization include the reduction of semi-infinite convex programming to finite subproblems [BTRBI79] or approximation algorithms [LS09].

In *topology*, Helly numbers come in two flavors. On the one hand, one may derive Helly numbers from topological conditions. The first result in this direction is *Helly's topological theorem* [Hel30], which states that a family of homology cells in $\mathbb{R}^d$ such that every subfamily of at most d members intersect in a homology cell has Helly number at most $d+1$. This was extended and generalized in several directions, for instance to take into account properties such as whether the sets separate the space or not [Had63]. Several results generalized Helly's topological theorem by allowing the elements of the families to have (and intersect into) several connected components [Mat97, KM08]; we will come back later to this line of research as we present a new result of this flavour. On the other hand, Helly numbers say something about the intersection patterns of families of sets, and these intersection patterns are classically studied via the *nerve* simplicial complex. In that setting, the Helly number bounds the maximum dimension of an induced simplicial subcomplex isomorphic to the boundary of a simplex. This naturally situates Helly's topological theorem as a particular case of more general results such as Borsuk's *Nerve theorem* [Bor48, Bjö03] or Leray's *acyclic cover theorem* [BT82] that relate the homotopy type or homology of a nerve of a family to that of its union. It also relates Helly number to other indicators of simplicial complexes such as the *Leray number* [KM08], the *representability* or the *collapsibility* [MT08], or the size of *blockers* in sparse representations of simplicial complexes [ALS11].

And the list goes on. In algebra, Helly numbers relate to the combinatorics of generators for certain (algebraic) groups [Far09]. In commutative algebra, they arise in the resolution of square-free monomial ideals [KM06] and, via their generalizations, multigraded ideals [Flø11]. We stop the enumeration here, as if the reader hasn't already been convinced that Helly numbers are natural, useful objects then we fear there is little else we can do to win him over.

⋆ ⋆ ⋆

Before we discuss our results some terminology is in order. Given a subset X of $\mathbb{R}^d$ we let $\mathcal{T}(X)$ denote the set of lines intersecting X. If $\mathcal{F} = \{X_1, \ldots, X_n\}$ is a family of subsets of $\mathbb{R}^d$, we let $\mathcal{T}(\mathcal{F}) = \bigcap_{X \in \mathcal{F}} \mathcal{T}(X)$ denote the set of lines intersecting every member of $\mathcal{F}$, the so-called *line transversals* to $\mathcal{F}$, and let $\mathcal{F}^{\mathcal{T}}$ denote the family $\mathcal{F}^{\mathcal{T}} = \{\mathcal{T}(X) \mid X \in \mathcal{F}\}$. We refer to the members of $\mathcal{F}$ as the *objects*. To save breath, we call the Helly number of $\mathcal{F}^{\mathcal{T}}$ the *transversal Helly number* of $\mathcal{F}$ and speak of a *family of disjoint sets* to designate a family whose members are pairwise disjoint.

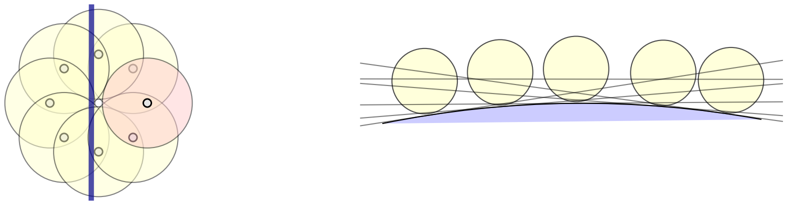

Figure 1: Two examples of constructions leading to arbitrary large transversal Helly numbers for convex sets. (Left): a family of n unit disks centered at the vertices of a regular n-gon. (Right): a family of n disks consisting of one large disk and $n-1$ small disks arranged so that any tangent line to the large disk misses at least one small disk.

As illustrated in Figure 1, there are families of convex sets in the plane with arbitrary large transversal Helly number. Such families can even consist of disjoint objects or of translates of the same object. Yet, families of *disjoint translates* of a convex set have bounded transversal Helly number. The first result in this direction was obtained by Ludwig Danzer.

Theorem 1.2 (Danzer [Dan57])**.** *The transversal Helly number of any family of disjoint unit disks is at most* 5*.*

The construction of Figure 1 (left) with $n = 5$ shows that this bound is best possible. Danzer's theorem was later extended by Grünbaum [Grü58] to families of disjoint translates of a square; Grünbaum then conjectured that the same bound holds for the transversal Helly number of any family of disjoint translates of a convex planar figure, a conjecture settled in the positive by Tverberg [Tve89] some four decades later.

Our starting points are two conjectures on generalizations of Danzer's theorem, not to other shapes, but to higher dimension. For $d \geqslant 2$, let $\mathcal{H}_d$ denote the maximum transversal Helly number of a family of pairwise disjoint unit balls in $\mathbb{R}^d$ (if there are such families with arbitrary large transversal Helly number we put $\mathcal{H}_d = \infty$). After proving that $\mathcal{H}_2 = 5$, Danzer conjectured:

Conjecture 1.3. *The number $\mathcal{H}_d$ is finite for any $d \geqslant 2$.*

It is easy to see that the number $\mathcal{H}_d$ are non-decreasing. Indeed, any family $\mathcal{F}$ of disjoint unit balls in $\mathbb{R}^d$ can be turned into a family $\tilde{\mathcal{F}}$ of disjoint unit balls in $\mathbb{R}^{d+1}$ with all centers in the hyperplane $x_{d+1} = 0$; since projecting a line orthogonally on

$x_{d+1} = 0$ decreases its distance to all balls' centers, $\mathcal{F}$ and $\tilde{\mathcal{F}}$ have the same transversal Helly number. Danzer made a second conjecture:

Conjecture 1.4. $\mathcal{H}_{d+1} > \mathcal{H}_d$ *for any* $d \geqslant 2$*.*

We will refer to Conjectures 1.3 and 1.4 as, respectively, Danzer's *upper bound conjecture* and *monotonicity conjecture*.

The first positive result on Danzer's upper bound conjecture was obtained in 1957 by Hadwiger [Had] for the case of families of "thinly distributed" balls; here, a family of balls is *thinly distributed* if the distance between any two balls' centers is at least twice the sum of their radii. This result was extended by Ambrus, Bezdek and Fodor [ABF06] to disjoint unit balls, in arbitrary dimension, the centers of which are distance at least $2\sqrt{2+\sqrt{2}}$ apart. Danzer's conjecture for three-dimensional disjoint unit balls, without additional assumption on their distribution, was only settled in 2001 by Holmsen, Katchalski and Lewis [HKL03]. Since $\mathcal{H}_2 = 5$ and the numbers $\mathcal{H}_d$ are non-decreasing, we have that $\mathcal{H}_d \geqslant 5$ for any $d \geqslant 2$; no better upper bound was known. In Part i, we prove that

$$2d - 1 \leqslant \mathcal{H}_d \leqslant 4d - 1, \tag{1.1}$$

settling the upper bound conjecture and providing evidence in support of the monotonicity conjecture.

The cornerstone of our proof is a convexity theorem for sets of directions of line transversals (Chapter 2). The space of directions in $\mathbb{R}^d$ is $\mathbb{S}^{d-1}$, envisaged as the unit sphere centered at the origin. The standard metric on $\mathbb{S}^k$ induces a notion of convexity, called *strong convexity*, on subsets $X \subset \mathbb{S}^k$ that do not contain any antipodal pair: X is strongly convex if it contains the smallest circle arc joining any two of its points. An oriented line transversal to a family $\mathcal{F}$ of disjoint balls induces an order on $\mathcal{F}$, namely the order in which the line meets the balls, that we call the *order induced* by the line on $\mathcal{F}$. We prove that if $\mathcal{F}$ is a family of disjoint balls in $\mathbb{R}^d$ then the set of directions of oriented line transversals to $\mathcal{F}$ that induce the same order is strongly convex. The special case of this convexity theorem was previously established for disjoint unit balls in $\mathbb{R}^3$ by Holmsen, Katchalski and Lewis [HKL03] using analytical methods. This proof was extended by Ambrus, Bezdek and Fodor [ABF06] to their setting[1] but has been observed to fail in the general case [GS05]. Our proof takes a different path, and relies on a careful inspection

1 Hadwiger [Had] also used a particular case of this convexity property for thinly distributed families of balls. That proof was apparently never published, and we do not know what arguments he used.

of the algebraic curves that compose the boundary of this set of directions.

The convexity theorem allows us to bound a local analogue of the Helly number (Chapter 3). Define the *pinning number* of a family $\mathcal{C}$ of subsets of some topological space as the largest integer k such that there exists $x_1, \ldots, x_k$ in $\mathcal{C}$ and $p \in \bigcap_{1 \leqslant i \leqslant k} x_i$ satisfying two conditions:

(I) p is an isolated point of $\bigcap_{1 \leqslant i \leqslant k} x_i$, and

(II) p is not an isolated point of $\bigcap_{1 \leqslant i \leqslant k; i \neq j} x_i$ for any j in $\{1, \ldots, k\}$.

We call a statement bounding from above a pinning number a *pinning theorem*. As usual, a point p is isolated in a set S if p is not a limit point of S. We show that if $\mathcal{F}$ is a finite family of disjoint balls in $\mathbb{R}^d$ the pinning number[2] of $\mathcal{F}^{\mathcal{T}}$ is at most $2d-1$. We further show that this bound is best possible by studying the stability of isolated transversal to a family of balls under tangency-preserving perturbations of the balls.

The last ingredient needed to obtain Inequalities 1.1 is a proof that the oriented line transversals to a family of $n \geqslant 9$ disjoint unit balls in $\mathbb{R}^d$ induce at most two pairs of reversed orders, which differ by the swapping of two consecutive elements (Chapter 4). From there, we obtain the upper bound by an homotopy argument: we start from a situation where every subset of $4d-1$ balls has a line transversal, deform the configuration until one of the $(4d-1)$-tuple has an isolated line transversal, and analyze that situation using our pinning theorem and the structure of the orders induced by line transversals. The lower bound on $\mathcal{H}_d$ follows easily from the lower bound on the pinning number.

To conclude this outline of our first series of results, let us mention that the study of transversal Helly numbers (for lines and higher dimensional transversals) opened a broader field of inquiry now known as *geometric transversal theory*. We refer the interested reader to the classical surveys of that area [DGK63, GPW93, Eck93, Wen04, GP02].

⋆ ⋆ ⋆

The pinning theorem (Theorem 3.5) that we use to prove the upper bound conjecture of Danzer holds for families of disjoint

2 The space of line is equipped with its natural topology, as defined e.g. through Plücker coordinates or via its identification with the quotient of the space of pairs of distinct points by the equivalence relation of "defining the same line".

balls without restriction on the radii, a situation in which the transversal Helly number cannot be bounded (see Figure 1 right). In Part ii, we explore other situations where pinning numbers of sets of line transversals can be bounded. This requires a new approach as our proof of Theorem 3.5 relies on the convexity structure of the sets of directions of line transversals to disjoint balls, which is quite specific to that setting.

Let us introduce some terminology. If a line transversal ℓ to a family $\mathcal{F}$ cannot move without missing some $A \in \mathcal{F}$, then we call the line ℓ *pinned* by $\mathcal{F}$. Here we consider small continuous movements of ℓ in the vicinity of its current position, obviously excluding translations parallel to itself. In other words, ℓ is pinned if ℓ is an isolated point in the space of line transversals to $\mathcal{F}$; we also call $\mathcal{F}$ a *pinning* of ℓ. If $\mathcal{F}$ pins ℓ but no proper subset of $\mathcal{F}$ does then we call $\mathcal{F}$ a *minimal pinning* of ℓ; the pinning number of $\mathcal{F}^{\mathcal{T}}$ is then simply the maximum size of a subset of $\mathcal{F}$ that is a minimal pinning.

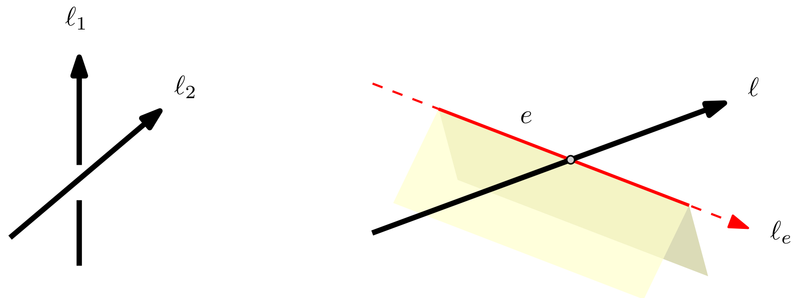

Figure 2: (Left): ℓ_2 passes to the right of ℓ_1. (Right): In the vicinity of ℓ, a line intersects the polytope if and only if it passes to the right of ℓ_e.

We first study the size of minimal pinnings of a line by polytopes in $\mathbb{R}^3$ (Chapter 5). We show that this pinning number can be arbitrarily large in general but can be bounded under a genericity condition: any minimal pinning of a line by polytopes in $\mathbb{R}^3$ where no polytope's facet is coplanar with the line has size at most eight. Given two non-parallel lines ℓ_1 and ℓ_2 with direction vectors $\vec{v_1}$ and $\vec{v_2}$, we say that ℓ_2 *passes to the right of* ℓ_1 if ℓ_2 can be translated by a positive multiple of $\vec{v_1} \times \vec{v_2}$ to meet ℓ_1; the genericity condition ensures that in the vicinity of the pinned line, intersecting a polytope is equivalent to a *conjunction* of sidedness constraints with respect to lines supporting that polytope's edges (See Figure 2). An adequate parameterization of line space as the points of a quadratic hypersurface $\mathfrak{M}$ in $\mathbb{R}^5$ recasts each sidedness constraint as a linear inequality. We thus have a polyhedral cone $C = \bigcap_{i \in I} H_i^+$, given as an intersection of halfspaces, whose apex a (the pinned line) is an isolated point of $C \cap \mathfrak{M}$, and we want a subset of at most eight of the halfspaces to define a (larger) cone C' such that a remains an isolated point of $C' \cap \mathfrak{M}$. We prove that such a subset indeed exists using a simple characterization of that isolation property in terms of the trace of the cone on the hyperplane tangent to the quadratic hypersurface in the apex a.

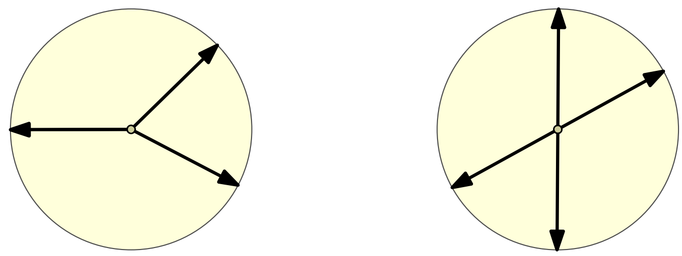

Figure 3: The two types of minimal families of vectors surrounding the origin in $\mathbb{R}^2$.

A natural follow-up question is whether the (generic) minimal pinnings of a line by polytopes in $\mathbb{R}^3$ can be tabulated. We give such a tabulation in the case where the sidedness constraints describing the conditions of intersecting the polytopes are induced by lines orthogonal to the pinned line (Chapter 6); we call such conditions *orthogonal constraints*. An adequate parameterization maps the set of lines passing to the right of an orthogonal constraint to a halfspace in $\mathbb{R}^4$. The minimal pinnings are thus recast as minimal families of halfspaces in $\mathbb{R}^4$ intersecting in a single point. Since each halfspace contains that point on its boundary, the question amounts to describing the situations where the outer normals of these halfspaces contain the origin in the interior of their convex hull; we thus speak of a family of normals that *surrounds* the origin (see Figure 3). We first characterize minimal families of vectors in $\mathbb{R}^4$ surrounding the origin; this characterization is given in terms of decomposition of the family in *critical simplices*, which are minimal linearly dependant subfamilies that contain the origin in their *relative* convex hull. We then identify those critical simplices that can be realized as normals to halfspaces defined by orthogonal constraints, and obtain a tabulation of all minimal pinnings of a line by orthogonal constraints (into 16 cases).

Our analysis of pinnings by constraints holds the key to a full classification of minimal stable pinnings by smooth convex sets, which we explore in Chapter 7. By *stable pinning* we mean a collection of objects that pins a line ℓ and that keeps pinning ℓ after small (independent) screws of axis ℓ are applied to each of the objects. We start by showing that, in an adequate representation of the space of lines, the set $\partial\mathcal{T}(C)$ of lines tangent to C is smooth in

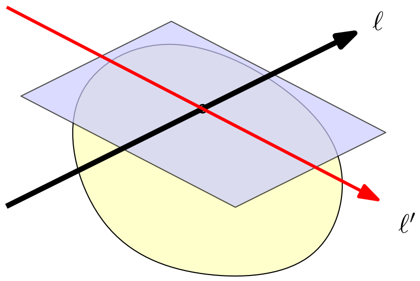

Figure 4: The first-order approximation in ℓ of the set of lines intersecting C is the set of lines passing to the right of ℓ'.

any line ℓ_0 that touches C in a smooth point of positive Gaussian curvature. Furthermore, the tangent space to $\partial\mathcal{T}(C)$ in ℓ_0 can be interpreted as the set of lines intersecting some orthogonal constraint (see Figure 4). We can thus study the situation where the first-order approximations of the solids $\mathcal{T}(C_i)$ intersect in a single point using techniques developped in Chapters 5 and 6. This condition of "pinning at first-order" turns out to be equivalent to the condition that the pinning be stable. The classification of minimal stable pinnings by smooth convex sets thus follows from the tabulation of minimal pinnings by orthogonal constraints. We also extend some arguments of Chapters 2 and 3 and establish a bound of 12 on the size of any minimal pinning of a line by convex sets in $\mathbb{R}^3$ under the condition that each set is tangent to the line in a smooth point of positive curvature and that no two sets are externally tangent on the pinned line. Contrary to the pinning theorem obtained in Chapter 3, all our arguments in Chapter 7 are *local*: we only need the sets to be convex, smooth, and have non-vanishing Gauss curvature in the vicinity of their contact points with the pinned line.

⋆ ⋆ ⋆

When $\mathcal{F}$ is a family of disjoint balls in $\mathbb{R}^d$ the intersections of members of $\mathcal{F}^{\mathcal{T}}$ are, in general, not connected. We know of two systematic ways to bound the Helly numbers of families of non-connected sets, and one may wonder if our upper bound on the transversal Helly number of families of disjoint unit balls could be amenable to such methods.

On the one hand, one can start with a "ground" family $\mathcal{G}$ whose Helly number is bounded and consider families $\mathcal{F}$ such that the intersection of any subfamily $\mathcal{G} \subseteq \mathcal{F}$ is a disjoint union of at most r elements of $\mathcal{H}$. When $\mathcal{G}$ is closed under intersection and *non-additive* in the sense that the union of disjoint elements of $\mathcal{G}$ is *never* an element of $\mathcal{G}$, the Helly number of $\mathcal{F}$ is at most r times the

Helly number of $\mathcal{G}$. This was conjectured (and proven for $r = 2$) by Grünbaum and Motzkin [GM61] and a proof of the general case was recently published by Eckhoff and Nischke [EN09], building on ideas of Morris [Mor73]. Direct proofs were also given by Amenta [Ame96] in the case where $\mathcal{G}$ is a finite family of compact convex sets in $\mathbb{R}^d$ and by Kalai and Meshulam [KM08] in the case where $\mathcal{G}$ is a good cover in $\mathbb{R}^d$ (recall that a *good cover* is a family where the intersection of any subfamily is empty or contractible).

On the other hand, it was shown by Matoušek [Mat97] and, independently, by Alon and Kalai [AK95] that if $\mathcal{F}$ is a family of sets in $\mathbb{R}^d$ such that the intersection of any subfamily is the union of at most r (possibly intersecting) convex sets, then the Helly number of $\mathcal{F}$ can be bounded from above by some function of r and d. Matoušek also gave a topological version of his theorem [Mat97, Theorem 2] where he bounds from above (again, by a function of r and d) the Helly number of families of sets in $\mathbb{R}^d$ assuming that the intersection of any subfamily has at most r connected components, each of which is $(\lceil d/2 \rceil - 1)$-connected, that is, has its ith homotopy group vanishing for $i \leqslant \lceil d/2 \rceil - 1$.

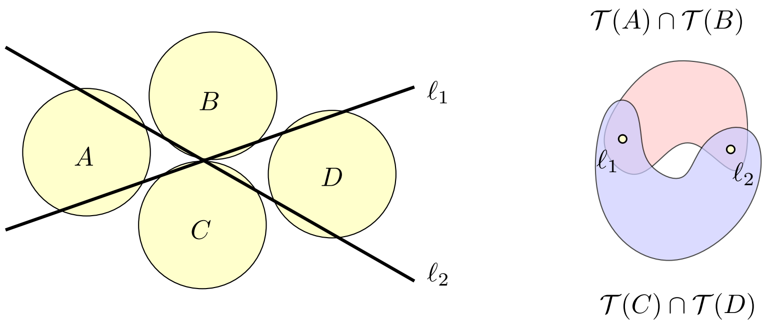

Figure 5: Four unit disks whose sets of transversals do not induce a good cover.

Our proof of the upper bound on the transversal Helly number of disjoint unit balls proceeds by arguing that the set of line transversals to any family of disjoint balls in $\mathbb{R}^d$ has a bounded number of connected components, each contractible.[3] Matoušek's theorem thus yields that these transversal Helly numbers are bounded, but provides no explicit upper bound.[4] The set of connected components of all intersections of subfamilies of $\mathcal{F}^{\mathcal{T}}$ do not, however, form a good cover (see Figure 5). The theorem

3 The contractibility fails for the case of *one* ball, as its set of line transversals has the homotopy type of $\mathbb{RP}^{d-1}$. This is, however, merely a technicality that can be taken care of by cutting the space adequately.

4 Matoušek's proof yields an explicit upper bound, but one so large that it is probably not interesting to nail it down precisely.

of Kalai and Meshulam, or, more generally, bounds of the first type therefore appear to be ineffective for our problem.

In Part iii we obtain a common generalization of the theorems of Matoušek and Kalai-Meshulam: a family of subsets of $\mathbb{R}^d$ where the intersection of any subfamily has at most r connected components, each of which is an homology cell, has Helly number at most $r(d+1)$. Our result holds, in fact, in any locally arc-wise connected topological space, the constant d being replaced by the smallest integer δ such that every open subset of that space has trivial $\mathbb{Q}$-homology in dimension δ and higher. We further allow some slack both on the topological condition (the connected components of $\bigcap_{A \in G} A$ may have nontrivial homology in low dimension) and the combinatorial condition (the bound r need only apply to the number of connected components of large enough subfamilies). [5] As a result, we sharpen inequality 1.1 into:

$$\mathcal{H}_d \leqslant 4d - 2 \quad \forall d \geqslant 6.$$

We also derive, in a similar way, a bound of 10 on the transversal Helly number of any family of disjoint translates of a convex planar figure, obtaining a weaker form of Tverberg's transversal theorem.

We obtain our general condition by following the approach developed by Kalai and Meshulam [KM08]. The intersection patterns of a family $\mathcal{F}$ of sets can be analyzed via the simplicial complex formed by its subfamilies with non-empty intersection, also called its *nerve*. The natural projection from the set C of connected components of members of $\mathcal{F}$ onto $\mathcal{F}$ itself extends into a simplicial map from the nerve of C to the nerve of $\mathcal{F}$. The key idea in the approach of Kalai and Meshulam is that the Helly number of $\mathcal{F}$ can be controlled via the homology of its nerve. Theorems *à la* Grünbaum-Motzkin simply follow from understanding how homology behaves under projection. In general, a projection can create homology in arbitrary high dimension; the Grünbaum-Motzkin setup ensures that the projection has fibers with bounded cardinality, which makes all the difference.

We make this general idea clearer by reformulating the (geometric) proof of Amenta's theorem and the (combinatorial) proof of the Eckhoff-Morris-Nischke theorem in this setting (Chapter 8). We then present the proof of our general condition which is based on similar ideas applied to a new refinement of the nerve complex which we introduce, the *multinerve* (Chapter 9). In a nutshell, the multinerve is a simplicial poset that encodes the

5 Even this relaxation is not, *stricto sensu*, a generalization of Matoušek's result as he allows nontrivial *homotopy* in *high* dimension whereas we allow nontrivial *homology* in *low* dimension. We ignore this technicality in first approximation.

intersection pattern of a family of subsets of a topological space. Where the nerve associates a simplex to every subfamily with non-empty intersection, the multinerve associate a simplex to every *connected component* of the intersection of every subfamily. We show that the classical *Nerve theorem* extends to multinerves and that the projection theorem underlying Kalai and Meshulam's proof extends to projection from the multinerve of $\mathcal{F}$ onto its nerve.

⋆ ⋆ ⋆

NOTES

Each chapter ends with a NOTES section, similar to this one, that gathers remarks on how the content of that chapter relates to other chapters or previous works. To help the reader identify the contributions, I use a numerical system, such as [1], for publications I co-authored and an alphabetical system, such as [Hel23], for other references.

For the ease of the presentation, I replaced some long or technical arguments by more high-level outlines. Most statements come with a reference to some published work; when developements posterior to the publication allowed for generalization at little cost it is reflected in the manuscript's statement.

This text assumes a general background in various areas of mathematics. For the convenience of the reader with a background on computational or combinatorial geometry, I included two appendices on, respectively, models of line space and simplicial homology.

Part I

DANZER'S CONJECTURES

2

CONE OF DIRECTIONS OF DISJOINT BALLS

The *cone of directions* of a family $\mathcal{F}$ of subsets of $\mathbb{R}^d$ is the subset $\mathcal{K}(\mathcal{F}) \subset \mathbb{S}^{d-1}$ of directions of oriented line transversals to $\mathcal{F}$ (see Figure 6). The main result of this chapter is a convexity theorem for cones of directions of families of disjoint balls.

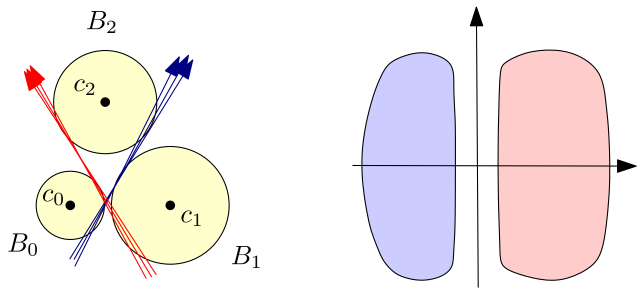

Figure 6: A triple T of balls with line transversals in two orders (left) and a planar representation of $\mathcal{K}(T)$ (right).

Recall that a subset of $\mathbb{S}^d$ is *strongly convex* if it contains no two antipodal points and contains the smallest circle arc joining any two of its points. A *strongly, strictly convex* set is a strongly convex set whose boundary contains no arc of great circle. We prove the following:

Theorem 2.1 (Convexity Theorem [8, Theorem 1])**.** *The directions of oriented line transversals to a family of at least two, and finitely many, disjoint balls in $\mathbb{R}^d$ in a given order form a strongly, strictly convex subset of $\mathbb{S}^{d-1}$.*

Theorem 2.1 therefore implies that for the cone of directions $\mathcal{K}(\mathcal{F})$ of a family $\mathcal{F}$ of disjoint balls is a disjoint union of strongly, strictly convex regions, one per order induced on $\mathcal{F}$ by its line transversals (see Figure 6). This convexity property will be instrumental in the next two chapters; we refer to Theorem 2.1 as our *convexity theorem*.

NOTATIONS. Most of this chapter is concerned with the case of three balls in $\mathbb{R}^3$, from which the general case follows quite easily. We let $T = \{B_0, B_1, B_2\}$ denote a triple of disjoint balls in $\mathbb{R}^3$, and let c_i and r_i denote, respectively, the center and radius of B_i. We call the triangle $c_0c_1c_2$ the *triangle of centers*. We assume the balls to be in generic position as Theorem 2.1 for *generic* triples T will extend, by limiting arguments, to special cases (Lemma 2.7). In particular we assume that the triangle of centers is non-degenerate, and call the plane spanned by c_0, c_1 and c_2 the *plane of centers*.

⋆ ⋆ ⋆

Our first step is to characterize the directions that appear on the boundary of $\mathcal{K}(T)$. Let u be a direction in $\mathbb{S}^2$ and let P_u denote the plane passing through the origin and orthogonal to u. The function that maps a line with direction v to its intersection point with P_v identifies the set of line transversals to a set X with the orthogonal projections of X on P_v. This leads to the following simple geometric description of $\partial\mathcal{K}(T)$.

Lemma 2.2 (Lemma 9 [13])**.** *If $u \in \mathbb{S}^2$ is a direction and I_u denotes the intersection of the orthogonal projections of B_0, B_1, B_2 on P_u, then:*

(i) $u \in \partial\mathcal{K}(T)$ *if and only if* I_u *is a single point, and*

(ii) $u \in Int(\mathcal{K}(T))$ *if and only if* I_u *has non-empty interior.*

Proof. The orthogonal projection of B_i on P_u is a disk of radius r_i whose center depends continuously on u. Let I_u denote the intersection of these three disks. Clearly, T has a line transversal with direction u if and only if I_u is non-empty. A non-empty intersection of disks has non-empty interior or is a single point. If I_u has non-empty interior then there is a neighborhood N of u such that for any $v \in N$, I_v is non-empty. It follows that any direction u such that I_u has non-empty interior belongs to the interior of $\mathcal{K}(T)$.

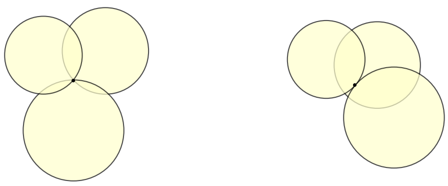

Figure 7: Two situations where three disks intersect in a single point.

When three disks intersect in a single point $\{x\}$, either x belongs to the boundary of all three disks or x is interior to one disk and a point of external tangency of the two other disks (see Figure 7). In each case, we can exhibit (see Figure 8) arbitrarily small perturbations v of u such that the intersection of the projections of B_0, B_1, B_2 on P_v is empty, showing that u must be on the boundary of $\mathcal{K}(T)$. □

By Lemma 2.2, the boundary of $\mathcal{K}(T)$ consists of directions of common *tritangents* to the three balls and of directions of *bitangents* to two of the balls contained in a plane tangent to and

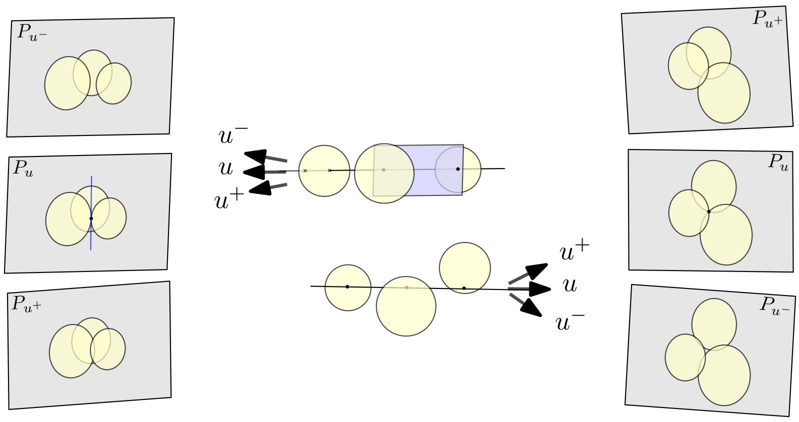

Figure 8: Perturbations of directions of projections u where I_u is a single point: in each case, u^+ is in the interior of $\mathcal{K}_{\prec}(T)$ and u^- is outside $\mathcal{K}_{\prec}(T)$.

separating these balls; such bitangents must pass through the inner center of similitude of the two balls and are therefore called *inner special bitangents*.

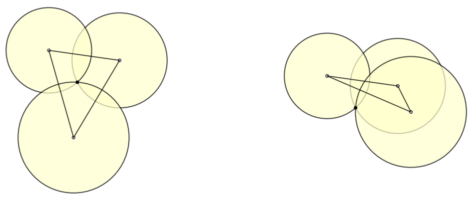

Figure 9: The direction of a tritangent may be on the boundary (left) or in the interior of $\mathcal{K}(T)$.

A direction of inner special bitangent that is a line transversal to the three balls is always on the boundary of $\mathcal{K}(T)$. A direction of tritangent may, however, be on the boundary or in the interior of $\mathcal{K}(T)$ (see Figure 9). The following simple geometric condition discriminates between these two situations.

Lemma 2.3 (Proposition 3 [8])**.** *The direction of a tritangent ℓ to B_0, B_1, B_2 belongs to $\partial\mathcal{K}(T)$ if and only if ℓ intersects the triangle of centers.*

Proof. If ℓ does not intersect the triangle of centers, we can translate it so that all three distances from ℓ to the balls' centers decrease. It follows that some translate of ℓ intersects the three open balls and the direction of ℓ belongs to $\mathrm{Int}(\mathcal{K}(T))$ by Lemma 2.2. Assume, on the other hand, that ℓ intersects the triangle of centers and let u denote the direction of ℓ. We consider the projection $c_0'c_1'c_2'$ of the triangle of centers on P_u and let $x = \ell \cap P_u$. If x could be moved to decrease the distance to all three centers, the sum of the areas of the three triangles $xc_i'c_j'$ would decrease. Since the areas of these three triangles sum up to the area of $c_0'c_1'c_2'$, this is impossible and no translation of ℓ intersect the interior of the three balls. It follows that the direction of ℓ belongs to $\partial\mathcal{K}(T)$ by Lemma 2.2. □

⋆ ⋆ ⋆

The boundary of $\mathcal{K}(T)$ consists of directions of inner special bitangents to pairs of balls and directions of tritangents to the three balls. We now consider these sets as algebraic curves and argue that in any point on the boundary of $\mathcal{K}(T)$ the corresponding curve is *locally* convex. It is technically more convenient to study these curves in the projective plane by working with unoriented direction in $\mathbb{RP}^2$. Since we are only interested in *local* convexity, this switch to a projective viewpoint has no consequence.

The directions of inner special bitangents to two disjoint balls make up a conic. The directions of common tangents to B_0, B_1, B_2 make up an algebraic curve of degree six in $\mathbb{RP}^2$. One way to see this is to begin with a description of lines in $\mathbb{R}^3$ by parameters $(p, u) \in \mathbb{R}^3 \times \mathbb{RP}^2$, where p is the orthogonal projection of the origin on the given line, and u is the direction of the line. We thus have

$$p \cdot u = 0,$$

and, after translating the triple of balls so that B_0 is centered at the origin, the conditions that the line with parameters (p, u) be tangent to B_0, B_1 and B_2 are equivalent to [6, Lemma 1]:

$$\begin{aligned}
\|p\|^2 &= r_0^2,\\
c_1 \cdot p &= \frac{\|c_1 \times u\|^2}{2\|u\|^2} + \frac{r_0^2 - r_1^2}{2},\\
c_2 \cdot p &= \frac{\|c_2 \times u\|^2}{2\|u\|^2} + \frac{r_0^2 - r_2^2}{2}.
\end{aligned}$$

Now, eliminating p from the above system yields a condition on u alone, characterizing the *direction-sextic*. Letting $e_{ij} = c_j - c_i$ and

$$t_{ij} = t_{ji} = \|e_{ij} \times u\|^2 = \|e_{ij}\|^2\|u\|^2 - \left(e_{ij} \cdot u\right)^2,$$

the result can be given by means of the Cayley determinant:

$$\sigma(u) = \det \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & \|u\|^2 r_0^2 & \|u\|^2 r_1^2 & \|u\|^2 r_2^2 \\ 1 & \|u\|^2 r_0^2 & 0 & t_{01} & t_{02} \\ 1 & \|u\|^2 r_1^2 & t_{01} & 0 & t_{12} \\ 1 & \|u\|^2 r_2^2 & t_{02} & t_{12} & 0 \end{pmatrix} = 0.$$

The fact that the elimination allows the stated Cayley determinant expression is given a natural explanation in [Bor06], but can be directly verified by computation.

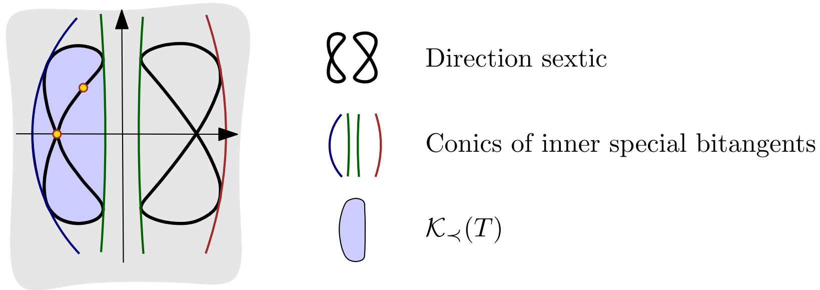

Figure 10: A representation of $\mathcal{K}_\prec(T)$ and the curves composing it. Observe that the direction sextic contain both singularities and flexes.

The sources of non-convexity in a real plane algebraic curve are two-fold: *singularities*, where the curve self-intersects or forms a cusp point, or *flexes*, where the curvature vanishes and may change sign. Both types of points can exist on the direction sextic (see Figure 10) but, as we now prove, only on arcs that are in the interior of the cone $\mathcal{K}(T)$. The flexes and singularities of an algebraic curve τ are the intersection of that curve with its Hessian curve [BK86]. Recall that the *Hessian* of a polynomial Q is defined as the determinant of the matrix of second derivatives:

$$H(Q) = H(Q)(u) = \det\left(\frac{\partial^2 Q}{\partial u_i \partial u_j}\right).$$

The Hessian curve of Q, or simply "the Hessian of Q", is the projective curve defined by the zero-set of this determinant.

Lemma 2.4 (Proposition 5 [8])**.** *An arc of the direction-sextic σ that belongs to the boundary $\partial\mathcal{K}(T)$ contains no flex or singularity of σ between its endpoints.*

Proof. Since the Hessian $H(\sigma)$ is an algebraic curve of degree twelve, the intersection of σ and $H(\sigma)$ has, counting multiplicities, $6 \times 12 = 72$ points, which leaves little hope for the possibility of "tracking" all flexes. Instead, following Lemma 2.3, we only consider directions of tangents to the three balls that cross the triangle of centers and are not directions of inner special bitangents. When projecting along such a tangent on a perpendicular plane, the projected centers $\tilde{c}_0 = 0, \tilde{c}_1, \tilde{c}_2$ form a triangle containing the point p, image of the tangent, as an interior point with squared distances r_i^2 to $\tilde{c}_i$. We equip $\mathbb{R}^3$ with a coordinate frame such that the projected triangle lies in the plane $e_3^\perp \subset \mathbb{R}^3$ and use three real parameters, x_0, x_1 and x_2, to describe the possible positions of the three centers:

$$c_0 = \tilde{c}_0 + x_0 e_3, \quad c_1 = \tilde{c}_1 + x_1 e_3, \quad c_2 = \tilde{c}_2 + x_2 e_3.$$

We express the corresponding direction-sextic σ and its Hessian $H(\sigma)$ as functions of $x = (x_0, x_1, x_2) \in \mathbb{R}^3$ depending on $\tilde{c}_0, \tilde{c}_1, \tilde{c}_2, r_0^2, r_1^2, r_2^2$. Now, all we have to prove is that

$$H(\sigma)(0,0,1) \neq 0$$

holds for all initial data (triangle and interior point) and all (x_0, x_1, x_2) corresponding to disjoint balls. In other words, we reduced the search for flexes on the *exterior* arcs of the direction sextic to the study of a polynomial function of x (and parameters). This function can be explicitly computed [8, Sections 4.2 and 4.3] and does, indeed, not vanish under the condition that the balls be pairwise disjoint. □

⋆ ⋆ ⋆

We now let $\prec$ denote some order on T, for instance $B_0 \prec B_1 \prec B_2$, and let $\mathcal{K}_\prec(T)$ denote the set of directions of oriented line transversals to $\mathcal{F}$ that induce $\prec$. We first observe that $\mathcal{K}_\prec(T)$ has a simple topological/homotopical structure.

Lemma 2.5 ([8, Proposition 4])**.** *If $\mathcal{K}_\prec(T)$ is non-empty then it is contractible. If $\mathcal{K}_\prec(T)$ is not reduced to a single point then it is the closure of its interior.*

Proof. Let ℓ be a line transversal to T in the order $\prec$. Let $\tilde{\ell}$ denote the reflection of ℓ with respect to the plane of centers. All lines between ℓ and $\tilde{\ell}$ (cf Figure 11) intersect the interior of the three balls and have therefore directions in $\mathrm{Int}(\mathcal{K}_\prec(T))$ by Lemma 2.2. We can thus retract $\mathcal{K}_\prec(T)$ onto the set A of directions of transversals to T (in the order $\prec$) in the plane of centers; since the balls are disjoint, A is a segment and $\mathcal{K}_\prec(T)$ is therefore contractible. If $\mathcal{K}_\prec(T)$ is not reduced to a single point, simple perturbations (see Figure 8) show that it is the closure of its interior. □

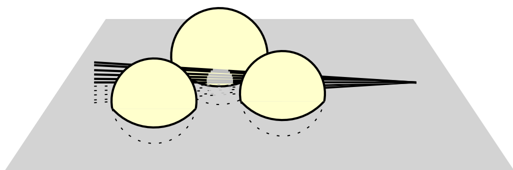

Figure 11: A component of line transversals to three balls can be retracted onto the set of transversals contained in the plane of centers.

We can now prove Theorem 2.1 for generic triples of disjoint balls.

Lemma 2.6 ([8, Lemme 9])**.** *If* T *is a triple of disjoint balls in generic position in* $\mathbb{R}^3$ *then* $\mathcal{K}_{\prec}(T)$ *is strongly, strictly convex.*

Proof. If $\partial\mathcal{K}_{\prec}(T)$ is made only of directions of inner special bitangents, strict convexity is immediate, since it is the intersection of strictly strongly convex regions bounded by conic arcs. Otherwise, the assumption that the balls are in generic position ensures that the direction-sextic σ is non-singular at all its contacts with any of the three conics determined by inner special tangents. Since the direction-sextic necessarily lies on the simply-connected side of each of the three conics, these contacts are tangency points at which $\partial\mathcal{K}_{\prec}(T)$ is locally convex. Thus, if we start at some point of $\partial\mathcal{K}_{\prec}(T)$ and follow the boundary curve, we obtain, by Lemma 2.4, a differentiable simple loop of class C^1, which is, locally, always on the same side of its tangent. Since, by Lemma 2.5, $\mathcal{K}_{\prec}(T)$ is contractible, its boundary consists of only one such curve. Now, for any affine plane $\mathbb{R}^2 \subset \mathbb{P}^2$ covering the loop, and any Euclidean metric in it, this means positive curvature on all its algebraic arcs and this implies [Topo6] that our simple loop bounds a compact convex set. In fact *strictly convex*, because of non-vanishing curvature. This proves the statement in the generic situation. □

The genericity assumption is taken care of by the following limit argument.

Lemma 2.7 ([8, Lemme 10])**.** *If* $\mathcal{K}_{\prec}(T)$ *has non-empty interior and* T *is the limit of a sequence of configurations* $T^{(\nu)}$ *such that* $\mathcal{K}_{\prec}(T^{\nu})$ *is strictly convex for each* ν *then* $K_{\prec}(T)$ *is strictly convex as well.*

Proof. Since, by Lemma 2.5, $\mathcal{K}_{\prec}(T)$ is the closure of its interior, to deduce the general case from the generic case it is enough to prove that, for any two points in the interior of $\mathcal{K}_{\prec}(T)$, the (geodesic) segment joining them is also contained in $K_{\prec}(T)$. Take two interior points. By assumption, for sufficiently large ν, the segment joining them is contained in all corresponding cones for $T^{(\nu)}$. Consider one point of the segment, and project the sphere configuration along the direction defined by the point, on a perpendicular plane. We have to prove that the disks representing the projected balls have at least one point in common. Suppose they don't. Then so would discs with the same centers and radii increased by a small $\epsilon > 0$. But then we can find, for sufficiently large ν, configurations $T^{(\nu)}$ with centers projecting less than $\epsilon/2$ away from those of T and corresponding radii with less than $\epsilon/2$ augmentation. Then the point of the segment cannot be in the respective cones of directions, a contradiction. Note that strict convexity still follows from non-zero curvature on smooth arcs for non-collinear centers, while for collinear centers it is obvious because of rotational symmetry. □

The general result now follows easily from the case of three balls in $\mathbb{R}^3$.

Theorem 2.1 (Convexity Theorem [8, Theorem 1]). *The directions of oriented line transversals to a family of at least two, and finitely many, disjoint balls in* $\mathbb{R}^d$ *in a given order form a strongly, strictly convex subset of* $\mathbb{S}^{d-1}$.

Proof. Recall that, for any family $\mathcal{F}$ of balls in $\mathbb{R}^3$, a direction will be realized by some transversal to $\mathcal{F}$ if and only if the orthogonal projection of the balls on a perpendicular plane has non-empty intersection. By Helly's Theorem in the plane, the cone of directions for a family of $n \geqslant 3$ balls is the intersection of the cones of directions of all its triples. Thus, the cone of directions of n ordered 3-dimensional disjoint balls is strictly strongly convex for any n.

Let $\mathcal{F}$ be a family of n disjoint balls in $\mathbb{R}^d$, let $\prec$ be an order on $\mathcal{F}$. Let u and v be two directions in $\mathcal{K}_\prec(\mathcal{F})$. We consider two line transversals ℓ_u and ℓ_v to $\mathcal{F}$ with respective directions u and v, and let E denote the 3-dimensional affine space these two lines span (or a 3-space containing their planar span, should the lines be coplanar). Now, E intersects $\mathcal{F}$ in a collection of 3-dimensional disjoint balls $\mathcal{F}'$ and any direction on the small arc of great circle joining u and v belongs to the space of direction of E. The statement for $d = 3$ implies that $\mathcal{K}_\prec(\mathcal{F}')$ is strictly strongly convex, and thus that for any any direction in the interior of the small arc of great circle joining u and v there exists a line intersecting the open balls of $\mathcal{F}'$ in the order $\prec$. These lines are also transversals to the interior of the balls in $\mathcal{F}$, and it follows that $\mathcal{K}_\prec(\mathcal{F})$ is strictly strongly convex. □

⋆ ⋆ ⋆

NOTES

This chapter is based on two articles co-authored with, respectively, Otfried Cheong, Andreas Holmsen and Sylvain Petitjean [11, 13] and Ciprian Borcea and Sylvain Petitjean [7, 8].

The convexity of sets of directions of line transversals was first investigated by Vincensini [Vin35]. He claimed that if $\mathcal{F}$ is a family of connected sets in the plane then the set $\mathcal{K}(\mathcal{F})$ of *non-oriented* directions of line transversals to $\mathcal{F}$ is convex. Vincensini considered directions in $\mathbb{RP}^1$ and called "convex" a subset that is the image of a strongly convex subset of $\mathbb{S}^1$ through the identification $\mathbb{S}^1/\mathbb{Z}_2 = \mathbb{RP}^1$. Vincensini's asserted that an intersection of "convex" sets in $\mathbb{RP}^1$ is again "convex", which is false (consider the green and orange sets in Figure 12). He then deduced that the Helly number of "convex" sets in $\mathbb{RP}^1$ is at most three. This is, again, false as can be seen by the following construction (illustrated in Figure 12 with $n = 5$). Start with n lines through the origin and associate to each line a "convex" set missing the two sectors incident to that line; any $n - 1$ sets meet a common line, but the intersection of all n of them is empty, so the Helly number of that family is n.

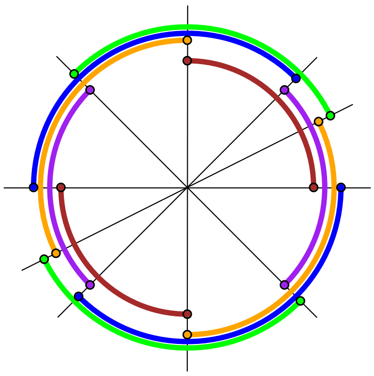

Figure 12: The Helly number of "convex" sets in $\mathbb{RP}^1$ is not bounded.

The convexity of sets of directions of line transversals was known to hold under three types of conditions:

- The convexity theorem was asserted by Hadwiger [Had] for *thinly distributed* families of balls, where the distance between any two balls' centers is at least twice the sum of their radii. This restriction ensures that there is only one order in which a line may intersect the balls, which simplifies the analysis. I am not aware of any published version of Hadwiger's proof and do not know which type of arguments he used.

- The convexity theorem was proven for the case of disjoint unit balls in $\mathbb{R}^3$ by Holmsen, Katchalski and Lewis [HKL03]. Their proof starts with two directions $u, v \in \mathcal{K}_{\prec}(\mathcal{F})$, construct a subset $Q \subset \mathbb{R} \times \mathbb{S}^1$ that project onto the intersection of $\mathcal{K}_{\prec}(\mathcal{F})$ by the great circle through u and v, and shows using analytical arguments that Q is convex.

- The proof technique of Holmsen et al. was applied by Ambrus, Bezdek and Fodor [ABF06] to families of unit balls in $\mathbb{R}^d$ such that the distance between any two centers is at least $2\sqrt{2+\sqrt{2}}$. What makes this distance condition interesting

> is that it is preserved under intersection by a subspace: if $\mathcal{F}$ satisfies the condition and E is a k-dimensional flat intersecting every ball in $\mathcal{F}$ then $E \cap \mathcal{F}$, as a collection of k-dimensional balls, also satisfies the condition. This allows to extend the property from three to arbitrary dimension.

The analytical approach of Holmsen et al. can be shown to fail for the case of disjoint balls of arbitrary radius in $\mathbb{R}^3$: the subset of $\mathbb{R} \times \mathbb{S}^1$ whose projection yields the intersection of $\mathcal{K}_{\prec}(\mathcal{F})$ by a great circle may not be convex. A different approach such as the one presented in this chapter was thus necessary.

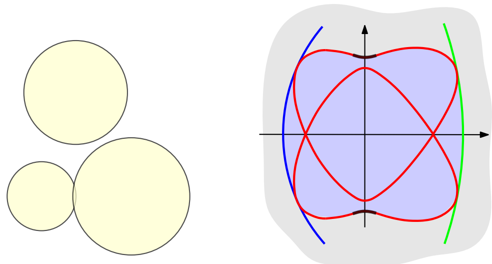

Figure 13: The strict convexity locally depends on the balls being disjoint globally.

If u is a direction of the boundary of $\mathcal{K}_{\prec}(\mathcal{F})$ then there is a unique line ℓ transversal to $\mathcal{F}$ with direction u. One could hope that the *local* convexity of the boundary of $\mathcal{K}_{\prec}(\mathcal{F})$ in u depends only on the disjointedness of the balls of $\mathcal{F}$ near ℓ. This is, unfortunately, not the case. Figure 13 shows three balls, two of which intersect, and the corresponding $\mathcal{K}_{\prec}(\mathcal{F})$. The non-convex portion of the boundary of $\mathcal{K}_{\prec}(\mathcal{F})$, marked in black, has positive length. The oriented lines tangents to the two intersecting balls in a common point make up a $\mathbb{S}^1$, at most 4 of which are also tangent to the third ball; this implies that there are only finitely many directions on the boundary of $\mathcal{K}_{\prec}(\mathcal{F})$ whose associated line does not meet the ball in distinct points.

Let us conclude this chapter by highlighting that the convexity theorem is rather specific to the case of balls. A natural way to measure how round a convex set is is to measure its *excentricity* (also called *fatness*): the ratio of the radius of the largest ball contained in the set to the radius of the smallest enclosing ball of the set. The excentricity is between 0 for a segment, ray or line and 1 for a ball. It turns out that for any $\varepsilon > 0$ there exists a family of convex sets of excentricity at least $1 - \varepsilon$ in $\mathbb{R}^3$ whose line transversals in a given order have directions that make up a non-connected subset of $\mathbb{S}^2$.

3

PINNING THEOREM FOR DISJOINT BALLS

In this chapter we consider a family $\mathcal{F}$ of disjoint balls in $\mathbb{R}^d$ (with arbitrary radii) and examine the pinning number of $\mathcal{F}^T$. We prove that this pinning number is always less than or equal to $2d-1$ and construct a family of examples showing that this bound is best possible for all $d \geqslant 2$.

NOTATIONS. Throughout this chapter, $\mathcal{K}(\mathcal{F}) \subset \mathbb{S}^{d-1}$ denotes the cone of directions of the family $\mathcal{F}$. If $\prec$ is an order on $\mathcal{F}$, we also use $\mathcal{K}_{\prec}(\mathcal{F})$ to denote the subset of $\mathcal{K}(\mathcal{F})$ of directions of line transversals to $\mathcal{F}$ realizing $\prec$.

⋆ ⋆ ⋆

By the Convexity Theorem, the projection, in the space of directions, of the sets of line transversals to disjoint balls is fairly simple, as it consists of a disjoint union of convex regions. We first observe that this simplicity can, to some extent, be pulled back to the sets of lines themselves.

Lemma 3.1 ([13, Lemma 14]). *If $\mathcal{F}$ is a family of $n \geqslant 2$ disjoint balls in $\mathbb{R}^d$ and D is a contractible subset of $\mathcal{K}(\mathcal{F})$ then the set of line transversals to $\mathcal{F}$ with direction in D is contractible.*

Proof. To any direction $u \in D$ we associate the line $\phi(u)$ with direction u and passing through the barycenter of the intersection of the orthogonal projection of the balls of $\mathcal{F}$ on the hyperplane through the origin with normal u (See Figure 14).

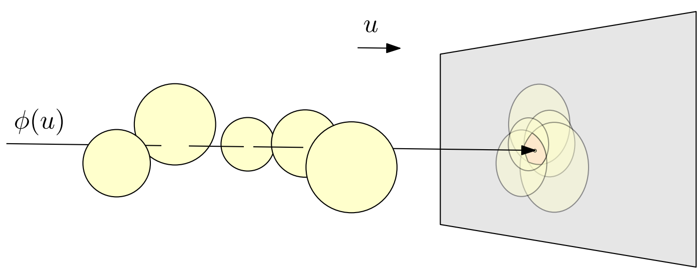

Figure 14: The barycentric transversal $\phi(u)$ in direction u.

Since $D \subset \mathcal{K}(\mathcal{F})$ that intersection is non-empty and $\phi(u)$ is thus well-defined (and is a line transversal to $\mathcal{F}$). It is clear that ϕ is one-to-one and that ϕ^{-1} is continuous; since the projection of a ball changes continuously with the direction of projection, ϕ is also continuous and is, in fact, an homeomorphism from D to $\phi(D)$. Since D is contractible, so is $\phi(D)$. Now, let L denote the set of line transversals to $\mathcal{F}$ with direction in D. Letting u_ℓ denote the direction of a line ℓ, the map

$$\begin{cases} L \times [0,1] \to L \\ (\ell, t) \mapsto \ell + t(\phi(u_\ell) - \ell) \end{cases}$$

is continuous. The set L is therefore homotopic to $\phi(D)$, which is contractible, so L itself is contractible. □

Lemma 3.1 and the Convexity Theorem imply that the set of line transversals to $n \geqslant 2$ disjoint balls in a given order is contractible (the cone of direction of a single ball being all of $\mathbb{S}^{d-1}$ is not contractible). This does not mean, however, that $\mathcal{F}^{\mathcal{T}}$ is a good cover (see Figure 15), as the restriction on the ordering is important.

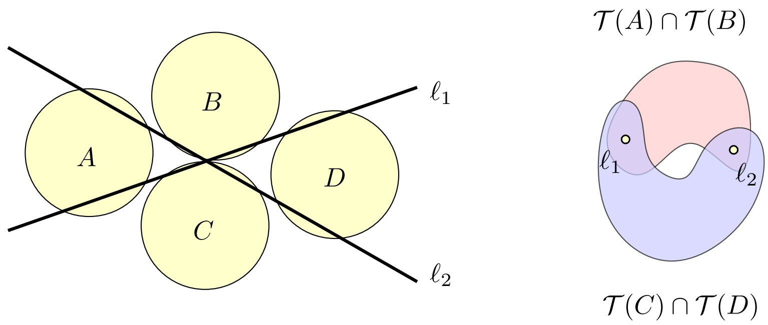

Figure 15: The set $\mathcal{F}^{\mathcal{T}}$ is not a good cover in general.

We can, however, restrict the space of lines to a subset U so that the family $\{\mathcal{T}(X) \cap U \mid X \in \mathcal{F}\}$ forms a good cover.

For any two distinct balls A and B in $\mathcal{F}$ we pick some hyperplane that separates them strictly. The directions parallel to this hyperplane form a hypersphere that splits $\mathbb{S}^{d-1}$ into two open hemispheres: $H(A,B)$, the set of directions pointing into the halfspace containing B and $H(B,A)$, the set of directions pointing into the halfspace containing A. To any order $\prec$ on $\mathcal{F}$ we associate:

$$R(\prec) = \bigcap_{A,B \in \mathcal{F} \mid A \prec B} H(A,B).$$

The set $R(\prec)$ is, when non-empty, an open, strongly convex subset of $\mathbb{S}^{d-1}$. In particular, $R(\prec)$ is non-empty for any order $\prec$ realized by an oriented line transversal to $\mathcal{F}$.

We now fix an order $\prec$ on $\mathcal{F}$ and let $\mathcal{L}$ denote the space of lines with direction in $H(A,B)$, where A and B is some fixed pair of balls in $\mathcal{F}$ with $A \prec B$. Let $\phi : \mathcal{L} \to (\mathbb{R}^{d-1})^2$ map each line of $\mathcal{L}$ to its intersections with two fixed translates of the hyperplane used to separate A and B. We let $\mathcal{L}_{\prec}$ denote the space of lines with direction in $R(\prec)$ and put

$$\mathcal{T}_{\prec}(B) = \phi(\mathcal{T}(B) \cap \mathcal{L}_{\prec}) \quad \text{and} \quad \mathcal{F}^{\mathcal{T}}_{\prec} = \{\mathcal{T}_{\prec}(B) \mid B \in \mathcal{F}\}.$$

Corollary 3.2. *The family $\mathcal{F}^{\mathcal{T}}_{\prec}$ is a good cover in $\mathbb{R}^{2d-2}$.*

Proof. Let $\mathcal{G} \subseteq \mathcal{F}$ and let $I = \bigcap_{B \in \mathcal{G}} \mathcal{T}_{\prec}(B)$. The set I is the image, under the mapping from the space of lines to $\mathcal{L}_{\prec}$, of the set J of line transversals to $\mathcal{G}$ with direction in $R(\prec)$. The intersection $D = R(\prec) \cap \mathcal{K}(\mathcal{G})$ is convex as $\mathcal{K}(\mathcal{G})$ is either $\mathbb{S}^{d-1}$, if $\mathcal{G}$ is a singleton, or convex by the Convexity Theorem, if $\mathcal{G}$ has cardinality two or more. Lemma 3.1 therefore yields that J is contractible. Since no line in J is parallel to $\partial H(A, B)$, the map $\mathbb{R}\mathbb{G} \to \mathbb{R}^{2d-2}$ induces an homeomorphism from J to I, and I is therefore contractible. □

With Helly's Topological Theorem, Corollary 3.2 immediately yields an analogue of Hadwiger's transversal theorem [Had57].

Corollary 3.3 ([13, Lemma 15])**.** *A family $\mathcal{F}$ of disjoint balls has a line transversal realizing the order $\prec$ if and only if every sub-family $\mathcal{G} \subseteq \mathcal{F}$ of size at most $2d - 1$ has a line transversal with direction in $R(\prec)$ that realizes the order induced on $\mathcal{G}$ by $\prec$.*

Before we prove our pinning theorem we need another, last, technical lemma.

Lemma 3.4 ([26, Lemma 3])**.** *Let ℓ be a line transversal to a finite family $\mathcal{F}$ of disjoint closed balls in $\mathbb{R}^d$ that realizes the order $\prec$. The following statements are equivalent:*

(i) *$\mathcal{F}$ pins ℓ,*

(ii) *ℓ is the only line transversal to $\mathcal{F}$ that realizes $\prec$,*

(iii) *$\mathcal{K}_{\prec}(\mathcal{F}) = \{u\}$ where u is the direction of ℓ,*

(iv) *there is no line transversal to $\mathcal{F}$ that realizes $\prec$ and intersects the interior of every ball.*

Proof. The equivalence of statements (i) and (ii) follows from the connectedness of the set of line transversals to $\mathcal{F}$ realizing $\prec$ (Lemma 3.1). The equivalence of statements (ii) and (iii) essentially follows from the fact that if $\mathcal{K}_{\prec}(\mathcal{F}) = \{u\}$ then u is on the boundary of $\mathcal{K}_{\prec}(\mathcal{F})$ and there is a unique line transversal to $\mathcal{F}$ in that direction (Lemma 2.2). The strict, strong convexity of $\mathcal{K}_{\prec}(\mathcal{F})$ (Theorem 2.1) implies that $\mathcal{K}_{\prec}(\mathcal{F})$ is a single point or has non-empty interior; the directions in the interior of $\mathcal{K}_{\prec}(\mathcal{F})$ are exactly the directions of line transversals to $\mathcal{F}$ that realize $\prec$ and intersects the interior of every ball (Lemma 2.2), so statements (iii) and (iv) are equivalent. □

We can now prove our pinning theorem.

Theorem 3.5 ([13, Proposition 13])**.** *Any minimal pinning of a line by finitely many disjoint closed balls in* $\mathbb{R}^d$ *has size at most* $2d-1$*.*

Proof. Let $\mathcal{F}$ be a family of $n \geqslant 2d$ disjoint closed balls in $\mathbb{R}^d$ that pins a line ℓ. Let $\prec$ be the order on $\mathcal{F}$ realized by ℓ. We let $\widetilde{\mathcal{F}}$ denote the collection of interiors of the balls in $\mathcal{F}$ and also write $\prec$ for the order on $\widetilde{\mathcal{F}}$ corresponding to $\prec$. Since $\mathcal{F}$ pins ℓ there is no line transversal to $\widetilde{\mathcal{F}}$ that realizes $\prec$ by Lemma 3.4 (iv). This means that the family $\widetilde{\mathcal{F}}^{\mathcal{T}}_{\prec}$ has empty intersection. This family is a good cover in $\mathbb{R}^{2d-2}$, by Lemma 3.1, so Helly's topological theorem implies that some subfamily $\widetilde{G} \subsetneq \widetilde{F}$ of cardinality at most $2d-1$ is such that $\widetilde{\mathcal{G}}^{\mathcal{T}}_{\prec}$ has empty intersection. This means that $\widetilde{G}$ has no line transversal with direction in $R(\prec)$. Note, however, that ℓ is a line transversal to the subfamily $G \subset \mathcal{F}$ corresponding to $\widetilde{G}$ with direction in $R(\prec)$.

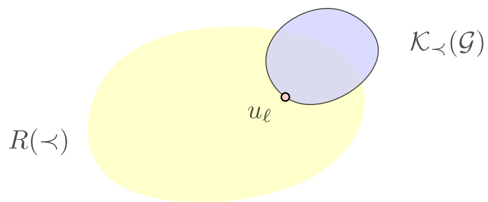

Figure 16: If $\mathcal{K}_{\prec}(\mathcal{G})$ has non-empty interior and intersects the open set $R(\prec)$ then the interior of $\mathcal{K}_{\prec}(\mathcal{G})$ also intersects $R(\prec)$.

In other words, $\mathcal{K}_{\prec}(\mathcal{G})$ intersects $R(\prec)$ but $\mathcal{K}_{\prec}(\widetilde{\mathcal{G}})$, which is the interior of $\mathcal{K}_{\prec}(\mathcal{G})$, does not intersect $R(\prec)$. Since $R(\prec)$ is open and $\mathcal{K}_{\prec}(\mathcal{G})$ is convex, it follows that $\mathcal{K}_{\prec}(\mathcal{G})$ must have empty interior (see Figure 16). Since $\mathcal{K}_{\prec}(\mathcal{G})$ is strictly convex, by the Convexity theorem, it is reduced to a single direction. We know that $\mathcal{K}_{\sigma}(\mathcal{G})$ contains the direction u_ℓ of ℓ, so $\mathcal{K}_{\sigma}(\mathcal{G}) = \{u_\ell\}$ and Lemma 3.4 (ii) implies that $\mathcal{G}$ pins ℓ. □

⋆ ⋆ ⋆

Let $\mathcal{F}$ be a collection of disjoint balls that pin a line ℓ and assume that each ball in $\mathcal{F}$ is tangent to ℓ. We call $\mathcal{F}$ a *stable pinning* of ℓ if there exists an $\varepsilon > 0$ such that ℓ is pinned by any family $\mathcal{F}' = \{B'_1, \ldots, B'_n\}$ of disjoint balls such that ball B'_i is tangent to ℓ and has its center at most distance ε away from the center of B_i; we call such a family $\mathcal{F}'$ an *ε-perturbation* of $\mathcal{F}$ (see Figure 17). If any sufficiently small perturbation of $\mathcal{F}$ minimally pins ℓ we call $\mathcal{F}$ a *stable minimal pinning* of ℓ. We now argue that in any dimension there exists a stable minimal pinning of a line by finitely many disjoint balls.

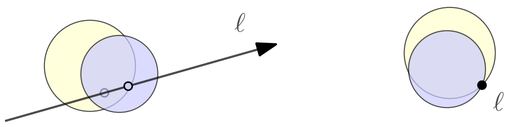

Figure 17: A (blue) perturbation of a (yellow) ball participating in the pinning of a line ℓ.

We start with a family $\mathcal{F} = \{B_1, \ldots, B_5\}$ of 5 disjoint balls tangent to a line ℓ in that order, and consider the projection of the balls on a plane perpendicular to ℓ. This projection consists of 5 disks with a common point on their boundary, the intersection point p of the line ℓ with the plane. Let $\vec{\eta}_i$ denotes the vector from the projection of the center of B_i to p. We say that the projection of $\mathcal{F}$ forms a σ_5 *pattern* if the cyclic order of the vectors $\pm\vec{\eta}_i$ is:

$$\vec{\eta}_1, -\vec{\eta}_3, \vec{\eta}_5, \vec{\eta}_2, -\vec{\eta}_4, -\vec{\eta}_1, \vec{\eta}_3, -\vec{\eta}_5, -\vec{\eta}_2, \vec{\eta}_4.$$

An example of geometric figure corresponding to a σ_5 pattern is depicted in Figure 18.

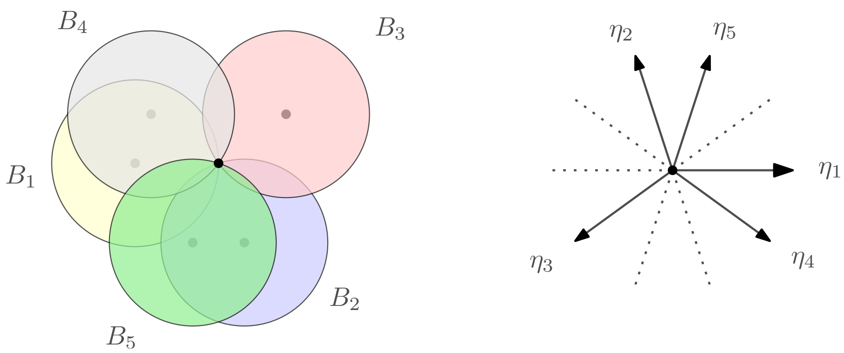

Figure 18: A σ_5-patterns.

Lemma 3.6 (Lemmas 1–4 [12])**.** *Let* $B_1, \ldots, B_5$ *be five disjoint balls in* $\mathbb{R}^3$ *tangent to a line* ℓ *in that order, and let* Π *be a plane orthogonal to* ℓ*. If the orthogonal projection of the* B_i *on* Π *is a* σ_5 *pattern then these five balls are a stable minimal pinning of* ℓ*.*

Proof. Let p denote the projection of the line ℓ, that is the point common to the projected disks. Let H_i denote the halfplane containing the projection of B_i and bounded by its tangent at p (see Figure 19). Let $\tilde{B}_i$ denote the interior of B_i. If ℓ is not pinned by the five balls, there exists, by Lemma 3.4 (iv), some line transversal ℓ' to $\tilde{B}_1, \ldots, \tilde{B}_5$ in that order. Since a small enough perturbation of ℓ' remains a transversal to the $\tilde{B}_i$, we can further assume that ℓ' is not coplanar with ℓ. This ensures that the orthogonal projection Δ of ℓ' on Π is a line that does not pass through p. It is easy (if a

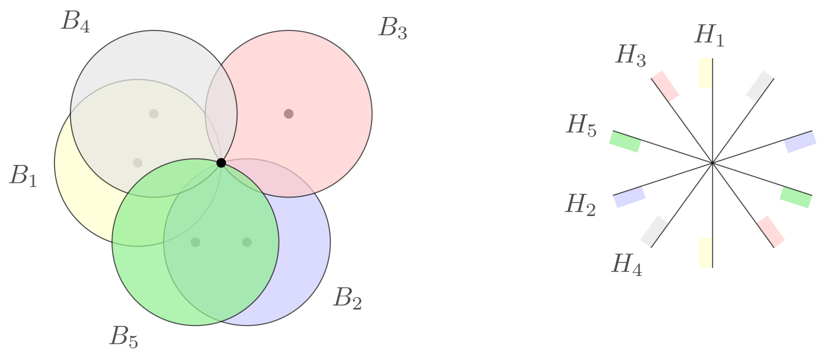

Figure 19: Turning the projected disk B_i (left) into the halfplane H_i (right).

bit tedious) to verify manually that there are always two indices $i < j$ such that Δ exits H_j before entering H_i; this contradicts the assumption that ℓ' exits B_i before entering B_j. Thus there is no such line ℓ' and Lemma 3.4 (iv) yields that the five balls pin ℓ.

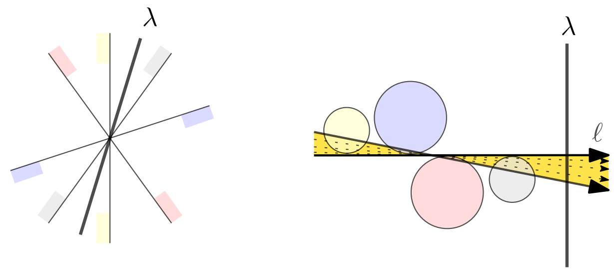

Figure 20: If a ball is removed (here B_5) we can find a plane inside which ℓ is not pinned.

If any ball is removed, it is, again, easy to check from the pattern that there must be a plane containing ℓ in which the traces of the balls do not pin ℓ, and minimality follows (see Figure 20). Since perturbing the balls keeps the projection a σ_5 pattern, it follows that the five balls form a stable minimal pinning of ℓ. $\square$

A family $\mathcal{F}$ of disjoint balls pin a line ℓ if and only if for every 3-flat E containing ℓ, the traces of the balls in E pin ℓ *inside* E. That the condition is necessary is obvious. That it is sufficient follows from Lemma 3.4 (ii): if ℓ is not pinned by $\mathcal{F}$ there exists some other line transversal ℓ' to $\mathcal{F}$ that realizes the same geometric permutation, and letting E denote a 3-flat containing[1] ℓ and ℓ', the existence of ℓ' implies, again by Lemma 3.4 (ii), that ℓ is not pinned in E by the traces of $\mathcal{F}$. With this property, and the compactness of the space of 3-flats containing a fixed line, we can use σ_5 patterns to build stable minimal pinnings in any dimension.

1 If ℓ and ℓ' are skew then E is unique.

Theorem 3.7 ([12, Theorem 4]). *For any $d \geqslant 2$, there exists a stable minimal pinning of a line by finitely many disjoint congruent balls in $\mathbb{R}^d$.*

Proof. Let ℓ be the x_d-axis in $\mathbb{R}^d$ and let Γ be the space of all three-dimensional flats containing ℓ. The natural homeomorphism between Γ and the space of two-dimensional linear subspaces of $\mathbb{R}^{d-1}$ implies that Γ is compact. For every $E \in \Gamma$ we can construct a quintuple Q_E of disjoint balls in $\mathbb{R}^d$ tangent to ℓ such that their restrictions to T project along ℓ to a σ_5-pattern. By construction, Q_E pins ℓ in E. By continuity, there exists a neighborhood N_E of E in Γ such that Q_E pins ℓ in any $E' \in N_E$. The union of all N_E covers Γ. Since Γ is compact, there exists a finite sub-family $\{E_1, \dots, E_n\}$ such that the union of the N_{E_i} covers Γ. Let $\mathcal{F}$ denote the union of the Q_{E_i}.

By construction, $\mathcal{F}$ is a finite family of balls such that the intersection of $\mathcal{F}$ with any 3-flat $E \in \Gamma$ is a stable pinning of ℓ in E. Let $\varepsilon > 0$ be such that any ε-perturbation $\mathcal{F}'$ of $\mathcal{F}$ remains a pinning of ℓ in each $E \in \Gamma$. Since a line is pinned by disjoint balls if and only if it is pinned in every 3-space that contains it, it follows that any ε-perturbation $\mathcal{F}'$ of $\mathcal{F}$ pins ℓ in $\mathbb{R}^d$, implying that $\mathcal{F}$ is a stable pinning. Since $\mathcal{F}$ is finite, some perturbation of one of its subfamily must be a minimal stable pinning of ℓ; let $\mathcal{G}$ denote that stable minimal pinning.

Now, we can perturb $\mathcal{G}$ so that it remains a stable minimal pinning and no two balls are tangent to ℓ at the same point. Observe that moving the center of a ball toward the contact point of that ball with ℓ, while reducing the radius to keep the ball tangent to ℓ, does not change any projection pattern. We can thus shrink the balls of $\mathcal{G}$, keeping the family a stable pinning of ℓ, until they are disjoint and have equal radius. □

⋆ ⋆ ⋆

We now argue that any stable minimal pinning must have cardinality at least $2d-1$. The proof is based on a simple interpretation of the first-order approximation of the set of lines intersecting a ball.

Consider, in $\mathbb{R}^d$, a line ℓ tangent to a ball B. Let p denote the point of contact of ℓ with B, H the hyperplane orthogonal to ℓ at p, H' the hyperplane tangent to B at p and H'_+ the closed halfspace bounded by H' containing B.

Definition 3.8. The *screen* of the pair (B, ℓ) is defined as the $(d-1)$-dimensional halfspace $H \cap H'_+$.

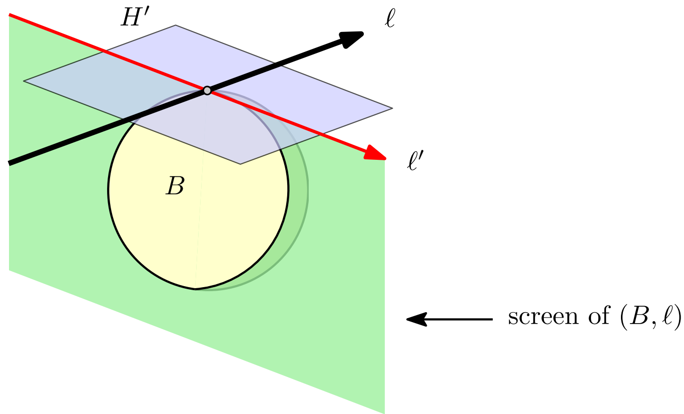

Figure 21: The screen of a ball and a line (in three dimension).

This definition is illustrated in Figure 21. Let $\mathcal{L}$ denote the space of oriented lines whose direction make a positive dot product with that of ℓ; as before, we let ϕ denote the parameterization of $\mathcal{L}$ by $\mathbb{R}^{2d-2}$ obtained by mapping each line of $\mathcal{L}$ to its intersections with two fixed hyperplanes perpendicular to ℓ. Let S be the screen of (B,ℓ) and denote by $\mathcal{T}(B)$ and $\mathcal{T}(S)$ the set of line transversals to B and S, respectively.

Lemma 3.9. *The set $\phi(\mathcal{T}(S))$ is a halfspace and the set $\phi(\mathcal{T}(B))$ is bounded by a degree 4 algebraic hypersurface, which is smooth at the origin. Moreover, $\mathcal{T}(B)$ and $\mathcal{T}(S)$ are internally tangent at the origin.*

Proof. Assume that ℓ is the x_d-axis and that we parameterize $\mathcal{L}$ using intersections with the hyperplanes $x_d = 0$ and $x_d = 1$.Let $\gamma = (u;v) \in \mathcal{L}$. Let α denote the x_d-coordinate of the point of tangency of ℓ and B and $\vec{\eta} = (\vec{\eta};0)$ the normal to B in its contact point with ℓ.

The line with parameters γ passes through $(u;0)$ and $(v;1)$ and therefore intersects the hyperplane $x_d = \alpha$ in the point $((1-\alpha)u + \alpha v;\alpha)$. This point belongs to S if and only if

$$\vec{\eta} \cdot ((1-\alpha)u + \alpha v) \leqslant 0.$$

Since

$$\begin{aligned} \vec{\eta} \cdot ((1-\alpha)u + \alpha v) &= ((1-\alpha)\vec{\eta}) \cdot u + (\alpha\vec{\eta}) \cdot v \\ &= ((1-\alpha)\vec{\eta};\alpha\vec{\eta}) \cdot (u;v), \end{aligned}$$

it follows that $\phi(\mathcal{T}(S))$ is the halfspace whose boundary contains the origin and with outer normal $((1-\alpha)\vec{\eta};\alpha\vec{\eta})$.

Let $\tilde{c} = (c;c_d)$ and r denote, respectively, the center and radius of B. The line with parameters γ intersects B if the equation

$$\|t(u;0) + (1-t)(v;1) - \tilde{c}\|^2 - r^2 = 0$$

has at least one real root. We thus have that $\mathcal{T}(B) = \{Q \geqslant 0\}$ where $Q(u;v)$ denotes the discriminant of this degree-two equation in t. Straightforward computations lead to

$$\begin{aligned} Q(u;v) = \quad & ((u-v)\cdot(v-c)+(c_d-1))^2 + r^2(\|u-v\|^2+1) \\ & -(\|u-v\|^2+1)(\|v-c\|^2+(1-c_d)^2) \end{aligned}$$

so $\phi(\mathcal{T}(B))$ is bounded by a degree 4 algebraic hypersurface whose gradient at the origin is $\overrightarrow{\text{grad}}Q_{(0;0)}(u;v) = ((1-c_d)c; c_d c)$. Since $c \neq 0$, as otherwise B would be centered on ℓ, it follows that $\phi(\mathcal{T}(B))$ is smooth at the origin. Since $\alpha = c_d$ and $\vec{\eta} = \lambda c$ for some $\lambda < 0$, it is immediate that $\phi(\mathcal{T}(B))$ and $\phi(\mathcal{T}(S))$ are internally tangent at the origin. $\square$

With Lemma 3.9, we can now argue that any family of $2d-2$ or fewer disjoint balls that pin a line admits arbitrarily close perturbations that do not pin that line. Since we know, by Lemma 3.7, that there exists a finite family $\mathcal{F}$ of disjoint balls in $\mathbb{R}^d$ that forms a stable pinning of a line, that configuration must have cardinality at least $2d-1$.

Theorem 3.10 ([12, Theorem 2]). *For any $d \geqslant 2$ there exists a family of $2d-1$ disjoint unit balls in $\mathbb{R}^d$ that minimally pins a line.*

Proof. Let $\mathcal{F} = \{B_1, \dots, B_n\}$ be a stable pinning of a line ℓ by disjoint balls in $\mathbb{R}^d$. We know that such a family exists by Lemma 3.7. Let S_i denote the screen of (B_i, ℓ) and given a perturbation $\mathcal{F}'$ of $\mathcal{F}$, we let B_i' and S_i' denote the corresponding perturbation of B_i and the screen of (B_i', ℓ). Assume that ℓ is the x_d-axis and let $\mathcal{L} \simeq \mathbb{R}^{2d-2}$ denote the space of parameters, in the stereographic coordinate system induced by $x_d = 0$ and $x_d = 1$, of all lines not orthogonal to ℓ. By Lemma 3.9, for $1 \leqslant i \leqslant n$ the solid $\mathcal{T}(B_i)$ is smooth in the origin and its first-order approximation at that point is $\mathcal{T}(S_i)$. Thus, any ray from the origin and contained, except for the origin, in the interior of $\mathcal{T}(S_i)$ must remain inside $\mathcal{T}(B_i)$ locally near the origin. It follows that if $\bigcap_{i \in J} \mathcal{T}(S_i)$ has non-empty interior for some $J \subset \{1, \dots, n\}$, the sub-family $\{B_i\}_{i \in J}$ cannot pin ℓ.

Consider a family of m or fewer halfspaces in $\mathbb{R}^m$ such that each halfspace has the origin on its boundary. It is a classical observation that if the intersection of these halfspaces has empty interior then the family of their outer normals is linearly dependent. As observed in the proof of Lemma 3.9, the outer normals to the $\mathcal{T}(S_i)$ live in the space

$$\{((1-\alpha)\vec{\eta}; \alpha\vec{\eta}) \mid \alpha \in \mathbb{R}, \vec{\eta} \in \mathbb{S}^{d-2}\}$$

which is nowhere locally contained in a hyperplane of $\mathcal{L}$ [12, Lemma 6]. Thus, for any sufficiently small $\varepsilon > 0$ there exists an

ε-perturbations $\mathcal{F}'$ of $\mathcal{F}$ with the property that the outer normals to any choice of $m \leqslant 2d-2$ halfspaces $\mathcal{T}(S_i')$ are linearly independent. Since $\mathcal{F}$ is stable, for ε small enough we know that $\mathcal{F}'$ pins ℓ and must therefore contain some minimal pinning $\mathcal{G}'$ of ℓ. Since no $m \leqslant 2d-2$ halfspaces $\mathcal{T}(S_i')$ have linearly dependant outer normals, $\mathcal{G}'$ must have cardinality at least $2d-1$; since any minimal pinning of a line by disjoint balls in $\mathbb{R}^d$ has cardinality at most $2d-1$ by Theorem 3.5, it follows that $\mathcal{G}$ consists in exactly $2d-1$ balls. □

⋆ ⋆ ⋆

NOTES.

This chapter is based on two articles co-authored with, respectively, Otfried Cheong, Andreas Holmsen and Sylvain Petitjean [11, 13] and Otfried Cheong and Andreas Holmsen [12].

The earliest appearance of the notion of pinned line that I am aware of is the use, in the proof of Hadwiger's transversal theorem [Had57], of the fact that if n disjoint convex sets in the plane pin a line then 3 of the sets suffices to pin the line.

Pinning is related to the concept of *grasping* used in robotics: an object is *immobilized*, or *grasped*, by a collection of contacts if it cannot move without intersecting (the interior of) one of the contacts (obviously, if the object is infinite, e.g. a line, one has to exclude those motions that leave it globally invariant). The study of objects ("fingers") that immobilize a given object has received considerable attention in the robotics community [MNP90, Mas01], and some Helly numbers for grasping are known. For instance, Mishra et al.'s work [MSS87] implies Helly-type theorems for objects in two and three dimensions that are grasped by point fingers. Theorem 3.5 can be interpreted in terms of grasping: whenever a cylinder is grasped by a family of fingers whose axis intersect the axis of the cylinder, a subset of at most $2d-1$ of these fingers always suffice to immobilized it; here the fingers could be either all outside or all inside the cylinder. Also, the notion of stable pinning is reminiscent of the notion of efficient grasping [KMY92].

Theorem 3.5 assumes that the balls are disjoint *globally* although the statement is about what happens *locally* near the pinned line. This is because our proof relies on the Convexity theorem; as noted at the end of Chapter 2, even if a line ℓ is pinned by a family $\mathcal{F}$ of balls that are disjoint near ℓ, it is possible that for some $\mathcal{G} \subseteq \mathcal{F}$ the set $\mathcal{K}_{\prec}(\mathcal{G})$ be non-convex in the direction of ℓ. We will explore in Part ii some *local* proofs of pinning theorems.

The orders induced by a line transversal to a family of sets have been extensively studied since the mid-1980's. If an order $\prec$ is realized by a line ℓ then the reverse orientation on ℓ realizes the order reverse of $\prec$; it is thus natural to consider *pairs* of orders, one reverse of another, usually dubbed *geometric permutations*. Separating hyperplanes such as the ones defining our $H(A, B)$ were used to obtain upper and lower bounds on the maximum number of geometric permutations of families of disjoint convex sets (or balls) in $\mathbb{R}^d$ [Wen90, SMS00, KV01, RKS10].

We obtain our pinning theorem by restricting the space of lines to a subset on which the family $\mathcal{F}^{\mathcal{T}}$ forms a good cover (Corollary 3.2) and applying Helly's topological theorem. This is, in some sense, similar to Grünbaum's argument for thinly distributed families of balls. Indeed, Grünbaum noted that the thin distribution condition ensures that in some well-chosen parameterization, the family $\mathcal{F}^{\mathcal{T}}$ *is* a good cover, and improved Hadwiger's quadratic bound by applying Helly's topological theorem. Some recent papers discuss (and fill) a gap in Helly's original proof of his topological theorem (see e.g. [KR06]); Helly's topological theorem for good covers does, however, follow easily from the classical nerve theorem [Bor48, Bjö03] via a connection that we explore in Chapters 8 and 9.

Let us conclude this chapter by a remark on Lemma 3.1, which observes that the space of line transversals to $\mathcal{F}$ in a given order has the same homotopy type as its projection onto the space of direction. Our proof proceeds by exhibiting a global section, but in general this is not necessary: the homotopy equivalence usually follows (assuming that the spaces considered are not too extravagant) from the simple fact that the projection has contractible fibers.

4

GEOMETRIC PERMUTATIONS & HELLY NUMBERS

In this chapter, we establish upper and lower bounds on the transversal Helly number $\mathcal{H}_d$ (defined on page 4). The proof of these bounds combine the Convexity theorem of Chapter 2 and the study of minimal pinnings of Chapter 3 with an analysis of the family of orders in which a family $\mathcal{F}$ of disjoint convex sets can be intersected by an oriented line.

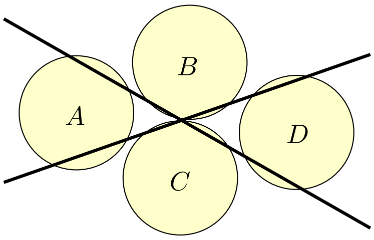

Figure 22: Two geometric permutations on four unit disks.

If a line ℓ induces the order $\prec$ on a family $\mathcal{F}$ then reversing the orientation of ℓ simply reverses the order $\prec$. It is therefore convenient to consider *geometric permutations*, that is pairs of orders, one reverse of the other (see Figure 22).

Theorem 4.1. *Let $\mathcal{F}$ be a family of n disjoint unit balls in $\mathbb{R}^d$. If $n \geqslant 9$ then $\mathcal{F}$ has at most two distinct geometric permutations that differ in the exchange of two adjacent elements. If $n \leqslant 8$ then $\mathcal{F}$ has at most three distinct geometric permutations.*

To prove Theorem 4.1, we show that certain pairs of geometric permutations cannot be simultaneously realized by line transversals to *four* disjoint unit balls, then bootstrap these restrictions by combinatorial arguments.

⋆ ⋆ ⋆

Let $\mathcal{F} = \{B_1, \ldots, B_n\}$ be a family of disjoint closed balls in $\mathbb{R}^d$. To discuss geometric permutations it is convenient to reduce to a situation where the geometric permutations are realized by isolated transversals. We do this using two types of transformations, which we call *shrinkings*.

The simplest transformation consists in decreasing the radius of the balls while *keeping their centers fixed*: we define $B_i(t)$ as the ball with same center as B_i and radius t times that of B_i and put $\mathcal{F}(t) = \{B_1(t), \ldots, B_n(t)\}$ (see Figure 23 left). If the centers of the balls in $\mathcal{F}$ are not all aligned then for any geometric permutation of τ there exists $t \in (0, 1)$ such that $\mathcal{F}(t)$ pins a line realizing τ.

Assume that we've shrunk $\mathcal{F}$ into $\mathcal{F}(t_1)$ while keeping the balls' centers fixed until it pins a line ℓ_1. We can further shrink $\mathcal{F}$ while *keeping the tangencies* by modifying the shrinking as follows: for $t < t_1$, if $B_i(t)$ does not intersect ℓ then we redefine it to be the ball with same radius and internally tangent to B_i in the projection of its center on ℓ (see Figure 23 right). Again, if $\mathcal{F}(t_1)$ has a geometric permutation τ_2 other than the one realized by ℓ_1 then $\mathcal{F}(t_2)$ pins both ℓ_1 and a line realizing τ_2 for some $t_2 \in (0, t_1]$.

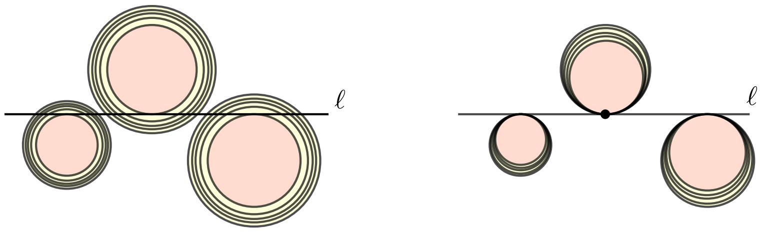

Figure 23: Shrinking while keeping the centers fixed (left) or while keeping tangencies to ℓ (right).

More generally, a *shrinking* of $\mathcal{F}$ is a 1-parameter family

$$(\mathcal{F}(t))_{t\in[0,1]} = (\{B_1(t), \ldots, B_n(t)\})_{t\in[0,1]}$$

such that (i) $\mathcal{F}(1) = \mathcal{F}$, (ii) the center and radii of each ball $B_i(t)$ are continuous functions of t, (iii) $B_i(t_2) \subset B_i(t_1)$ for any $1 \leqslant i \leqslant n$ and any $0 \leqslant t_1 < t_2 \leqslant 1$. The two previous claims follow from the next lemma.

Lemma 4.2. *Let $\mathcal{F}(t)$ be a shrinking of a family $\mathcal{F}$ of disjoint closed balls in $\mathbb{R}^d$. For every geometric permutation τ of $\mathcal{F}$ let I_τ denote the set of $t \in [0, 1]$ such that $\mathcal{F}(t)$ has a line transversal realizing τ. The set I_τ is a closed interval and if $\min I_\tau > 0$ then $\mathcal{F}(\min I_\tau)$ pins a line realizing τ.*

Proof. The inclusions $B_i(t_2) \subset B_i(t_1)$ for $0 \leqslant t_1 < t_2 \leqslant 1$ ensure that I_τ is an interval that contains 1. Let $(t_i)_{i\in\mathbb{N}}$ be a sequence converging to t^* such that $\mathcal{F}(t_i)$ has a line transversal ℓ_i in a given geometric permutation τ for all $i \in \mathbb{N}$. Since $\cup_{t\in[0,1]}\mathcal{T}(\mathcal{F}(t))$ is compact, there is a subsequence $(t_{i_j})_{j\in\mathbb{N}}$ such that ℓ_{i_j} admits a limit ℓ^*; since the balls are closed, this line ℓ^* is a line transversal to $\mathcal{F}(t^*)$ in the geometric permutation τ. It follows that I_τ is a closed interval of the form $[t^*, 1]$. Now, if $t^* > 0$ then $\mathcal{F}(t^*)$ has no line transversal realizing τ and intersecting the interior of every ball, as such a line would be a transversal to any $\mathcal{F}(t)$ with $t < t^*$ sufficiently close to t^*. It follows, with Lemma 3.4, that $\mathcal{F}(t^*)$ pins a line that realizes τ. □

⋆ ⋆ ⋆

We start with a geometric observation due to Otfried Cheong, Jae-Soon Ha and Jungwoo Yang [CHY10] which we call the *Distance lemma*.

Lemma 4.3 (Distance Lemma). *Let $B_1, \ldots, B_4$ be four disjoint unit balls in $\mathbb{R}^d$ and let c_i denote the center of B_i. If there exists an oriented line that intersects $B_1, \ldots, B_4$ in that order then $\|c_1c_4\|$ is strictly larger than $\|c_1c_2\|$, $\|c_2c_3\|$ and $\|c_3c_4\|$.*

Proof. It suffices to prove the statement for $d = 3$ as we can always project the line on a 3-flat containing the four centers and analyze the configuration in that space. Moreover, we can assume that the line is pinned by the four balls by first shrinking the balls while keeping their centers fixed until no line intersect the *interiors* of $B_1, \ldots, B_4$ in that order; the quadruple then has a unique line transversal in that order and we can make the balls' radii unit again by scaling that pinning configuration up.

Let Π be the plane perpendicular to ℓ and passing through the origin and let λ be the intersection point of ℓ with Π. Let t_i and p_i denote the projections of c_i on ℓ and Π, respectively. For $i = 1, 2, 3$ we let α_i denote the angle $\angle p_i \lambda p_{i+1}$, measured counterclockwise and note that $\angle p_1 \lambda p_4 = \alpha_1 + \alpha_2 + \alpha_3 \mod 2\pi$.

Let $g(x) = \sqrt{2 + 2\cos x}$. In the plane, the distance between two points of the unit circle that make a central angle of x is precisely $g(\pi - x)$, so $g(x+y) \leqslant g(x) + g(y)$. Since the balls are unit, $\|p_ip_{i+1}\| \leqslant g(\pi - \alpha_i)$, with equality if B_i and B_{i+1} are both tangent to ℓ. Also, from $\|c_ic_{i+1}\|^2 = \|t_it_{i+1}\|^2 + \|p_ip_{i+1}\|^2$, the balls' disjointedness and $g(\pi - x)^2 = 4 - g(x)^2$ we get that $\|t_it_{i+1}\| > g(\alpha_i)$.

If B_1 and B_4 are both tangent to ℓ we have:

$$\|c_1c_4\|^2 = (\|t_1t_2\| + \|t_2t_3\| + \|t_3t_4\|)^2 + 4 - g(\alpha_1 + \alpha_2 + \alpha_3)^2.$$

A simple computation shows that for $i \in \{1, 2, 3\}$ we have:

$$\begin{aligned}(\|t_1t_2\| + \|t_2t_3\| + \|t_3t_4\|)^2 > \quad & (g(\alpha_1) + g(\alpha_2) + g(\alpha_3))^2 \\ & + \|t_it_{i+1}\|^2 - g(\alpha_i)^2\end{aligned}$$

Combining the two previous relations we obtain

$$\|c_1c_4\|^2 > \Delta + \|t_it_{i+1}\|^2 + 4 - g(\alpha_i)^2,$$

where $\Delta = (g(\alpha_1)+g(\alpha_2)+g(\alpha_3))^2 - g(\alpha_1+\alpha_2+\alpha_3)^2 \geqslant 0$ and:

$$\|t_i t_{i+1}\|^2 + 4 - g(\alpha_i)^2 \geqslant \|t_i t_{i+1}\|^2 + \|p_i p_{i+1}\|^2 = \|c_i c_{i+1}\|^2.$$

Thus, $\|c_1 c_4\|$ is strictly larger than any of $\|c_1 c_2\|, \|c_2 c_3\|$ and $\|c_3 c_4\|$.

If B_1 and B_4 are not both tangent to ℓ we can assume, without loss of generality, that it is not tangent to B_4 and is therefore pinned by B_1, B_2, and B_3. We start with Pythagoras' identity:

$$\|c_1 c_4\|^2 - \|c_i c_{i+1}\|^2 = \|t_1 t_4\|^2 - \|t_i t_{i+1}\|^2 + \|p_1 p_4\|^2 - \|p_i p_{i+1}\|^2.$$

From $p_1 = p_3$ we have that $\|p_1 p_4\|^2 = \|p_3 p_4\|^2$ and $\|t_1 t_4\| > \|t_3 t_4\|$ implies that $\|c_1 c_4\| > \|c_3 c_4\|$. When $i = 1, 2$, $\|p_i p_{i+1}\| = 2$ and, with $\|p_1 p_4\|^2 = \|p_3 p_4\|^2 = \|c_3 c_4\|^2 - \|t_3 t_4\|^2$ we get:

$$\|c_1 c_4\|^2 - \|c_i c_{i+1}\|^2 = \underbrace{\|t_1 t_4\|^2 - \|t_i t_{i+1}\|^2 - \|t_3 t_4\|^2}_{\geqslant 0} + \underbrace{\|c_3 c_4\|^2 - 4}_{>0}.$$

Thus, $\|c_1 c_4\|$ is also larger than $\|c_1 c_2\|$ and $\|c_2 c_3\|$. □

We continue with a condition on the directions of line transversals to triples of disjoint unit balls which we call the *Angle lemma*.

Lemma 4.4 (Angle Lemma [15, Lemma 7]). *Let B_1, B_2, B_3 be disjoint unit balls in $\mathbb{R}^d$, $d \geqslant 3$, and let c_i denote the center of B_i. The direction of an oriented line that intersects B_1, B_2, B_3 in that order makes an angle of strictly less than $\pi/4$ with $\overrightarrow{c_1 c_3}$.*

Proof. The general case can be shown to follow from the three-dimensional case. Moreover, the Convexity theorem implies that it suffices to show that no oriented line transversal to B_1, B_2, B_3 in that order has a direction vector making angle exactly $\pi/4$ with $\overrightarrow{c_1 c_3}$. We can thus choose a coordinate system in which the line transversal ℓ and the line $(c_1 c_3)$ have respective direction vectors $-\vec{x} - \vec{y}$ and $-\vec{x}$. We translate B_1 by a positive multiple of $\vec{x}$ and B_3 by a negative multiple of $\vec{x}$ until they become tangent to ℓ, and observe that in the process all distances between centers increase. We assume, up to exchanging the roles of B_1 and B_3, that the x-coordinate of c_2 is non-negative and translate B_2 by a positive multiple of $-\vec{x} - \vec{y}$ until the x-coordinate of c_2 is 0; in the process the distance $\|c_1 c_2\|$ only increases. Elementary analytic considerations then show that the intersection of the plane $x = 0$ with the cylinder of axis ℓ and radius 1 is an ellipse contained in the intersection of the plane $x = 0$ with the ball centered in c_1 and radius 2. It follows that B_1 and B_2 intersect, a contradiction. □

A first consequence of the Distance Lemma is that it reduces the bounds on the number of geometric permutations stated in Theorem 4.1 to the case of four balls.

Lemma 4.5. *If* $n \geqslant 4$ *disjoint unit balls in* $\mathbb{R}^d$ *have three (resp. four) geometric permutations then three (resp. four) of the balls have three (resp. four) distinct geometric permutations.*

Proof. Let $\mathcal{F}$ be a family of n disjoint unit balls in $\mathbb{R}^d$. Call an element *extreme* in a geometric permutation if it appears first in one of its orders. We start with two observations:

(i) Any two geometric permutations of $\mathcal{F}$ have an extreme element in common. If $n \leqslant 3$ this is obvious. If $n \geqslant 4$, the union of two disjoint extreme pairs would be a quadruple with two geometric permutations for which the Distance Lemma yields contradicting inequalities.

(ii) If $n \geqslant 4$ and σ_1 and σ_2 are two geometric permutations with A as a common extreme elements then the restrictions of σ_1 and σ_2 to $\mathcal{F} \setminus \{A\}$ are distinct. Indeed, writing this restriction $BC \dots D$, we can assume that $\sigma_1 = ABC \dots D$ and $\sigma_2 = AD \dots CB$ and the Distance Lemma, applied to the restrictions of the σ_i to $\{A, B, C, D\}$, yields contradicting inequalities.

Assume that $\mathcal{F}$ has three geometric permutations and consider a minimal subfamily $G \subseteq \mathcal{F}$, with $|G| \geqslant 4$, on which these permutations are pairwise distinct. Any two geometric permutations must share an extreme element, by remark (i). If all three have an extreme element in common then remark (ii) implies that we can remove this element and keep the permutations distinct, contradicting the minimality of G. It follows that there exists a triple $\{A, B, C\}$ such that the three geometric permutations have $\{A, B\}$, $\{B, C\}$ and $\{A, C\}$ as extreme pairs, and the triple $\{A, B, C\}$ already distinguishes the three geometric permutations. This proves the first statement.

Assume that $\mathcal{F}$ has four geometric permutations $\sigma_1, \dots, \sigma_4$ and let $G \subseteq \mathcal{F}$, with $|G| \geqslant 5$, be a minimal subfamily on which these permutations are pairwise distinct. Let $T \subset G$ be a triple T on which the restrictions of $\sigma_1, \sigma_2, \sigma_3$ are pairwise distinct and assume, without loss of generality, that $\sigma_{4|G} = \sigma_{1|G}$. Orienting σ_1 and σ_4 so that elements of T appear in the same order, we find some pair $P \subset G$ on which σ_1 and σ_4 differ, and so all σ_i are distinct on $T \cup P$. We can thus assume that $|G| = 5$ and write $G = \{A, B, C, D, E\}$. The minimality of G and remark (ii) implies that there is no element extreme in all σ_i. Since any two σ_i have an extreme element in common, there exist three elements, say $\{A, B, C\}$, such that, say $\{A, B\}$, $\{B, C\}$ and $\{A, C\}$ are the extreme pairs of three of the geometric permutations, say σ_1, σ_2, and σ_3 respectively. Moreover, two elements from $\{A, B, C\}$ cannot be consecutive in σ_1, σ_2, or σ_3. Indeed, if, for instance, $\sigma_1 = ACDEB$ then the Distance Lemma yields that ab and ae are strictly larger than ac; then, writing $\sigma_3 = AXYZC$, one of X or Y must be B or E, and the Distance Lemma yields a contradicting inequality; the other cases are symmetric. The only possibility is thus

$$\sigma_1 = AX_1CY_1B, \quad \sigma_2 = BX_2AY_2C, \quad \sigma_1 = AX_3BY_3C,$$

where $\{X_i, Y_i\} = \{D, E\}$ and $X_i \neq Y_i$. Applying the Distance Lemma to quadruples of the σ_i we obtain

$$\|ay_1\| > \|ax_1\| \quad \text{and} \quad \|ay_3\| > \|ax_3\|, \quad \text{implying that } Y_1 = Y_3,$$
$$\|bx_1\| > \|by_1\| \quad \text{and} \quad \|by_2\| > \|bx_2\|, \quad \text{implying that } Y_2 = X_1,$$
$$\|cx_2\| > \|cy_2\| \quad \text{and} \quad \|cx_3\| > \|cy_3\|, \quad \text{implying that } Y_2 = Y_3,$$

and finally $Y_1 = Y_3 = Y_2 = X_1$, a contradiction. Thus, no quintuple minimally distinguishes four geometric permutations. □

⋆ ⋆ ⋆

Using the Distance and Angle lemmas, we can now show that for certain pairs (τ_1, τ_2) of geometric permutations there exists no family of disjoint unit balls in $\mathbb{R}^d$ with both a line transversal realizing τ_1 and one realizing τ_2. We call such pairs *incompatible* (for disjoint unit balls).

Proof of Theorem 4.1. The Distance Lemma yields contradicting inequalities when applied to the following pairs of geometric permutations on $\{1, 2, 3, 4\}$:

$$(1234, 4123), \quad (1234, 1432), \quad (1234, 3412), \quad (1234, 3142).$$

Naturally, if (σ, τ) are incompatible then so are $(\sigma \circ \upsilon, \tau \circ \upsilon)$ for any permutation υ. Figure 24 represents the pairs of geometric permutations for which imcompatibility follows from the Distance Lemma. It is easy to check that among any four geometric permutations on $\{1, 2, 3, 4\}$ there are two that are incompatible for disjoint unit balls. With Lemma 4.5, this proves that any family $\mathcal{F}$ of disjoint unit balls has at most three geometric permutations.

A simple packing argument [15, Lemma 6] shows that if the axis of two unit-radius cylinders make an angle of $\pi/4$ or more then the intersection of the solid cylinders cannot contain 9 points at pairwise distance at least 2. It follows that if $\mathcal{F}$ consists of at least 9 balls then the angle between any two line transversals to $\mathcal{F}$ is strictly less than $\pi/4$. With the Angle Lemma, this prevents $(1234, 1342)$ from arising as restrictions[1] of two geometric

1 We do not claim that $(1234, 1342)$ is an incompatible pair, merely that no two line transversals to *the whole family* can intersect four of the balls in the orders 1234 and 1342.

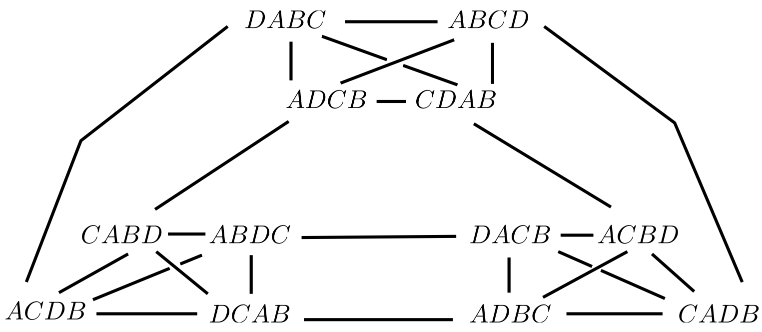

Figure 24: All geometric permutations on 4 elements. Full edges indicate an incompatible pair for disjoint unit balls.

permutations of $\mathcal{F}$ to four of its elements. Indeed, let σ and σ' be two line transversals to $\mathcal{F}$ that induce the geometric permutations $B_1B_2B_3B_4$ and $B_1B_3B_4B_2$, respectively, on a quadruple $\{B_1, B_2, B_3, B_4\} \subseteq \mathcal{F}$. We orient each of σ and σ' so that they meet 1 before 2, and let $\vec{n}$ and $\vec{n'}$ denote their respective direction vectors. Since σ and σ' meet 1 before 2, the angle between $\vec{n}$ and $\vec{n'}$ is at most $\pi/2$ and the above packing argument yields that the angle between $\vec{n}$ and $\vec{n'}$ is less than $\pi/4$. The Angle Lemma yields that $\vec{n'}$ makes an angle less than $\pi/4$ with the vector $\overrightarrow{c_3c_2}$, where c_i denotes the center of B_i. Thus, $\vec{n}$ must make an angle less than $\pi/4 + \pi/4 = \pi/2$ with $\overrightarrow{c_3c_2}$, implying that σ meets B_3 before B_2, a contradiction.

Figure 25 augments the diagram of Figure 24 by the pairs equal, up to relabelling, to $(1234, 1342)$. We can check that if $\mathcal{F}$ has 9 balls or more then for any quadruple $Q \subseteq \mathcal{F}$ the restrictions to Q of the geometric permutations of $\mathcal{F}$ consist of at most two geometric permutations of Q that differ by the swapping of two consecutive elements. With Lemma 4.5, this immediately implies the first statement of the theorem, and concludes this proof. $\square$

⋆ ⋆ ⋆

Recall that given a family $\mathcal{F} = \{C_1, \ldots, C_n\}$ of convex sets in $\mathbb{R}^d$ we let $\mathcal{F}^{\mathcal{T}} = \{T_1, \ldots, T_n\}$ where T_i denote the set of lines intersecting C_i. Recall that

$$\mathcal{H}_d = \max_{\mathcal{F} \in \mathrm{DUB}(d)} \text{Helly number of } \mathcal{F}^{\mathcal{T}},$$

where $\mathrm{DUB}(d)$ denotes the set of all finite families of disjoint unit balls in $\mathbb{R}^d$. Before we bound $\mathcal{H}_d$ we show that families of disjoint balls afford an analogue of Hadwiger's transversal theorem.

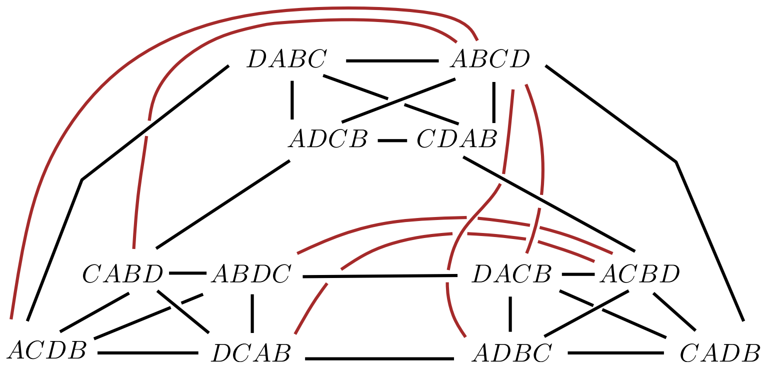

Figure 25: As Figure 24, with pairs equivalent, up to relabelling, to $(1234, 1342)$ connected by curvy edges.

Lemma 4.6. *Let $\mathcal{F}$ be a family of disjoint balls in $\mathbb{R}^d$ and let τ be a geometric permutation of $\mathcal{F}$. If every subset $G \subseteq \mathcal{F}$ of size $2d$ (resp. $2d+1$) has a line transversal in the geometric permutation $\tau_{|G}$ then $\mathcal{F}$ has a line transversal (resp. a line transversal in the geometric permutation τ).*

Proof. Let $(\mathcal{F}(t))_{t\in[0,1]}$ be a shrinking of $\mathcal{F}$ that keeps the centers fixed. Given $G \subseteq \mathcal{F}$ we denote by $G(t)$ the subfamily of $\mathcal{F}(t)$ corresponding to G. By Lemma 4.2, for any $G \subseteq \mathcal{F}$ of size $2d$ there is a $t_G \in [0,1]$ such that $G(t)$ has a line transversal realizing $\tau_{|G}$ for any $t \geqslant t_G$ and $G(t_G)$ pins that line. Let $t^* = \max_{G\in\binom{\mathcal{F}}{2d}} t_G$ and assume that $t^* = t_{G_0}$. Assume that $t^* > 0$ as otherwise the statement trivially holds. By Theorem 3.5, some subfamily $G_1(t^*) \subset G_0(t^*)$ of size at most $2d-1$ also pins ℓ. Let $B \in \mathcal{F}$ and let $H = G_1 \cup \{B\}$. By definition of t^*, $H(t^*)$ has a line transversal realizing $\tau_{|H}$. Since, by Lemma 3.4 (ii), ℓ is the only line transversal to G_1 that realizes $\tau_{|G_1}$, it follows that $B(t^*)$ intersects ℓ, and so does B. Since this holds for any $B \in \mathcal{F}$, the line ℓ is a line transversal to $\mathcal{F}$. This proves the first statement. The proof of the second statement is similar up to the definition of G_1. We then pick some $A \in G_1$ and observe, using the same arguments, that for any pair $\{B, B'\} \subset \mathcal{F}$ the triple $\{A, B, B'\}$ must meet ℓ consistently with τ. Since this holds for all pairs $B, B' \in \mathcal{F}$, it follows that ℓ is a line transversal to $\mathcal{F}$ realizing τ. □

We can now prove the main result of this part: upper and lower bounds on $\mathcal{H}_d$.

Theorem 4.7. *For any $d \geqslant 2$ we have $2d-1 \leqslant \mathcal{H}_d \leqslant 4d-1$.*

Proof. By Theorem 3.10, for any $d \geqslant 2$ there exists a family $\mathcal{F}$ of $2d-1$ disjoint unit balls in $\mathbb{R}^d$ that minimally pin a line ℓ. We shrink the balls while keeping them tangent to ℓ until they are small enough that they can be separated by hyperplanes orthogonal to ℓ. The resulting configuration is a minimal pinning of ℓ by $2d-1$ disjoint unit balls with no other line transversal. Every subfamily of size $2d-2$ admits a line transversal meeting the interior of the balls. Thus, reducing the radii of all balls by some $\varepsilon > 0$ small enough, while keeping the centers fixed, ensures that the $2d-1$ balls have no line transversal while every subfamily of $2d-2$ balls retains some line transversal. This shows that $\mathcal{H}_d \geqslant 2d-1$.

Now to the upper bound. Let $\mathcal{F}$ be a family of disjoint unit balls in $\mathbb{R}^d$ such that every subfamily of size $4d-1$ has a line transversal. Let $\mathcal{F}(t)$ be a shrinking of the balls that keeps the centers fixed. By Lemma 4.2, there is some t^* such that (i) for all subfamilies $G \subseteq \mathcal{F}$ of size $4d-1$, $G(t^*)$ has a line transversal, and (ii) for some subfamily $G_0 \subseteq \mathcal{F}$ of size $4d-1$, $G_0(t^*)$ pins all its line transversals. By Theorem 4.1 we are in one of two cases.

Case 1. $G_0(t^*)$ has two line transversals, ℓ_1 and ℓ_2, each realizing a distinct geometric permutation. By Theorem 3.5 for $i = 1, 2$ some subfamily $G_0^i \subseteq G_0$ of size at most $2d-1$ is such that $G_0^i(t^*)$ pins ℓ_i. Letting $G_1 = G_0^1 \cup G_0^2$ we thus have that $G_1(t^*)$ has size at most $4d-2$ and has exactly two line transversals ℓ_1 and ℓ_2. Any ball of $\mathcal{F}(t^*)$ thus intersects ℓ_1 or ℓ_2. If all balls intersect the same line then $\mathcal{F}$ has a line transversal and we are done. Otherwise we let A be some ball in $\mathcal{F}$ that misses ℓ_2; the subfamily $G_2 = G_1 \cup \{A\}$ has cardinality $4d-1$ and $G_2(t^*)$ has exactly one line transversal, which reduces to our second case.

Case 2. $G_0(t^*)$ has a single line transversal ℓ. By Theorem 3.5 some subfamily $G_1 \subset G_0$ of size at most $2d-1$ is such that $G_1(t^*)$ pins ℓ. For each $Z \in G_0 \setminus G_1$ let $G_Z = G_0 \setminus \{Z\}$. If one of the $G_Z(t^*)$ has no other line transversal than ℓ then the same reasoning as above yields that ℓ is a line transversal to $\mathcal{F}(t^*)$ and therefore to $\mathcal{F}$. We thus assume that every $G_Z(t^*)$ has a line transversal $\ell_Z \neq \ell$. By Theorem 4.1 the geometric permutations of G_Z realized by ℓ and ℓ_Z differ in the swapping of two consecutive elements X_Z and Y_Z. Since $G_1(t^*)$ pins ℓ, these geometric permutations must already differ on G_1 and for all Z we must have $\{X_Z, Y_Z\} \subseteq G_1$. For $Z_1, Z_2 \in G_0 \setminus G_1$ let $G_{Z_1,Z_2} = G_0 \setminus \{Z_1, Z_2\}$. The family $G_{Z_1,Z_2}(t^*)$ has at most two geometric permutations, by Theorem 4.1, and three transversals ℓ, ℓ_{Z_1} and ℓ_{Z_2}; it follows that ℓ_{Z_1} and ℓ_{Z_2} realize the same geometric permutation of G_{Z_1,Z_2} for any $Z_1, Z_2 \in G_0 \setminus G_1$, and the pair $\{X_Z, Y_Z\}$ is independent of Z; let us denote them by X and Y. Let τ be the geometric permutation of G_0 realized by ℓ and let τ' be the image of τ after swapping X and Y. Since $G_0(t^*)$ has no line transversal realizing τ' there exists, by Lemma 4.6, some subset $G_2 \subset G_0$ of size at most $2d+1$ such that $G_2(t^*)$ has no line transversal realizing $\tau'_{|G_2}$. Observe that G_2 must contain both X and Y as otherwise $\tau'_{|G_2} = \tau_{|G_2}$ and ℓ, which is a line transversal to $G_0(t^*) \supset G_2(t^*)$, would realize $\tau'_{|G_2}$. It follows that $G_3 = G_1 \cup G_2$ has size at most $4d-2$ and $G_3(t^*)$ has a unique line transversal ℓ. As a consequence, every ball in $\mathcal{F}(t^*)$ must intersect ℓ, and $\mathcal{F}$ has a line transversal. $\square$

⋆ ⋆ ⋆

NOTES.

This chapter is based on two articles co-authored with, respectively, Otfried Cheong and Hyeon-Suk Na [14, 15] and Otfried Cheong, Andreas Holmsen and Sylvain Petitjean [11, 13]; the proof of Theorem 4.1 [15] was remarkably simplified by the Distance Lemma of Otfried Cheong, Jae-Soon Ha and Jungwoo Yang [CHY10].

It was previously known that when n is sufficiently large the number of geometric permutations of n disjoint unit balls is at most 4, regardless of the dimension [HXC01, HXC04, KSZ03, ZS01]. Theorem 4.1 sharpens this previous bound while dropping the assumption that n be sufficiently large. The question whether families of $4 \leqslant n \leqslant 8$ disjoint unit balls in $\mathbb{R}^d$ can have three geometric permutations remains open. Theoretically, this could be solved mechanically: one could, for instance, compute a point in every connected component of the complement of the discriminant variety of some adequately defined parametric semi-algebraic set. In practice, this seems to remain out of reach of computer algebra systems as of 2011.

The previous bounds on the number of geometric permutations of n disjoint unit balls [HXC01, HXC04, KSZ03, ZS01] proceed by limiting the number of switch pairs[2] that can occur. Our proof of Theorem 4.1 uses the different approach of identifying *incompatible pairs* of patterns for geometric permutations of the same set of disjoint unit balls. Incompatible pairs were previously used by Asinowski [Asi99] to study geometric permutations of disjoint translates of a convex figure in the plane, and *incompatible families* of patterns for geometric permutations in higher dimension have been investigated by Asinowski and Katchalski [AK05].

Lemma 4.5 asserts that for disjoint unit balls, if any 4 balls have at most 3 geometric permutations then any family has at

2 A *switch pair* is a pair $\{X, Y\}$ such that there exist some other object Z such that the collection admits oriented line transversals meeting the triple in the orders $X \prec Y \prec Z$ and $Y \prec X \prec Z$, respectively.

most 3 geometric permutations. This invites the obvious question of looking at what happens if 4 and 3 are replaced by two parameters m and k. This is essentially the analogue for geometric permutations of classical questions on *shatter functions* and the theory of *VC-dimension*, also known as *Dirac-type problems* in extremal graph theory; we investigated this question in a separate work with Otfried Cheong and Cyril Nicaud [16].

Incompatible pairs are reminiscent of *excluded patterns* in (algebraic) combinatorics where, along the lines of the celebrated Stanley-Wilf conjecture (now the Marcus-Tardos theorem [MT04]) one studies the maximum number of permutations on n elements that do not have a fixed permutation on m elements as a restriction. Any permutation on n elements can arise as a geometric permutation of a family of n disjoint unit balls in $\mathbb{R}^d$, as one can always start with a line transversal and relabel the balls so that they appear along the line in the desired geometric permutation; it is therefore natural that conditions on geometric permutations involve *two or more* geometric permutations. Incompatible pairs for geometric permutations of course imply excluded patterns for standard permutations by assuming one of the permutations to be the identity, but that naive reduction loses information (as we ignore conditions between pairs not involving the identity geometric permutation). It would be interesting to further explore how the techniques designed for the study of excluded patterns could be applied to the geometric permutation setting.

Our proof that $\mathcal{H}_d \leqslant 4d-1$ uses the characterization of pairs of geometric permutations given by Theorem 4.1. As we will see in Chapters 8 and 9, the bound on the number of geometric permutations is, in principle, sufficient to obtain such a bound on the Helly number.

The smallest enclosing cylinder of a set of n points in the plane can be computed in $\Theta(n \log n)$ time [EW89, ARW89]. In higher dimension[3], this problem is more difficult: in three dimension, the best known algorithm has complexity $O(n^{3+\varepsilon})$ and if the dimension is part of the input then it is NP-hard [Meg90]. It was known that in the plane the problem is easier, and can be solved in $O(n)$ time, for *sparse* point sets, i.e. point sets where the radius of the smallest enclosing cylinder is smaller than twice the smallest inter-point distance [Gar92]. The upper bound on h_d implies, via the relation between Helly numbers and the combinatorial dimension of LP-type problems [Ame94], that in *any fixed dimension* d the smallest enclosing cylinder of a sparse point set can be computed in $O(n)$ time.

3 A cylinder in $\mathbb{R}^d$ is understood as the set of points within a given distance from a line.

Part II

PINNING THEOREMS IN $\mathbb{R}^3$

5 PINNING BY POLYTOPES

In this chapter, we prove the following pinning theorem for sets of line transversals to polytopes:

Theorem 5.1. *Any minimal pinning of a line by possibly intersecting convex polytopes in $\mathbb{R}^3$, no facet of which is coplanar with the line, has size at most eight.*

This statement is similar to Theorem 3.5, albeit with a genericity constraint: the pinned line is not allowed to lie in the plane of a polytope's facet. This condition may be stronger than is necessary, but *some* condition must remain as we show that there can be no pinning theorem for sets of line transversals to arbitrary polytopes:

Theorem 5.2. *There exist arbitrarily large minimal pinnings of a line by convex polytopes in $\mathbb{R}^3$.*

NOTATIONS. In this chapter, all lines are oriented unless specified otherwise and ℓ_0 denotes the line to be pinned. If ℓ is a line we denote by $\overrightarrow{\ell}$ its direction, in $\mathbb{S}^2$. Given two non-parallel lines ℓ_1 and ℓ_2 with direction vectors $\overrightarrow{\ell_1}$ and $\overrightarrow{\ell_2}$, we say that ℓ_2 *passes to the right of* ℓ_1 if ℓ_2 can be translated by a positive multiple of $\overrightarrow{\ell_1} \times \overrightarrow{\ell_2}$ to meet ℓ_1, or, equivalently, if

$$\det \begin{pmatrix} p_1 & p_1' & p_2 & p_2' \\ 1 & 1 & 1 & 1 \end{pmatrix} < 0, \tag{5.1}$$

where p_i and p_i' are points on ℓ_i such that $\overrightarrow{p_i p_i'}$ is a positive multiple of $\overrightarrow{\ell_i}$. We shorten $\bigcup_{x \in F} x$ and $\bigcap_{x \in F} x$ into, respectively, $\bigcup_F$ and $\bigcap_F$.

⋆ ⋆ ⋆

A line g meeting ℓ_0 in a single point represents a *constraint* on ℓ_0. A line ℓ *satisfies* a constraint g if and only if ℓ meets g or passes to the right of g.

A polytope whose interior intersects a line cannot contribute to a minimal pinning of that line. We can thus restrict our attention to polytopes tangent to ℓ_0. Let P be a polytope tangent to ℓ_0 such that no facet of P is coplanar with ℓ_0, and observe that ℓ_0 must

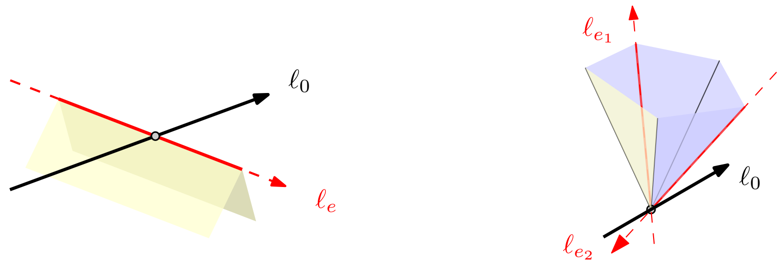

Figure 26: The line ℓ_0 may touch P in the interior of an edge (left) or in a vertex (right).

therefore intersect P in a single point p. We are in one of two cases. On the one hand, p may be interior to an edge e of P. Let Π_e denote the plane spanned by ℓ_0 and e. We let ℓ_e be the line supporting e, oriented so that $\overrightarrow{\ell_0} \times \overrightarrow{\ell_e}$ points in the halfspace bounded by Π_e and containing P (see Figure 26). A line ℓ near ℓ_0 intersects P if and only if it passes to the right of ℓ_e. On the other hand, p may be a vertex of P. In that case, the non-coplanarity condition implies that exactly two edges e_1 and e_2 are incident to p and span, with ℓ_0, a supporting plane of P. Let ℓ_{e_i} denote the line supporting e_i, oriented with the same rule as above. A line ℓ near ℓ_0 intersects P if and only if it passes to the right of *both* ℓ_{e_1} and ℓ_{e_2}.

Now consider a pinning of ℓ_0 by polytopes so that no facet of a polytope is coplanar with ℓ_0. Since pinning is a local property, we can ignore the polytopes and speak only about the lines supporting their relevant edges. We can thus consider pinnings of ℓ_0 by constraints; here ℓ_0 is *pinned by a family of constraints* if it is an isolated point in the space of lines satisfying those constraints. In particular, to prove Theorem 5.1 it suffices to show that any minimal pinning of ℓ_0 *by constraints* has size at most eight.

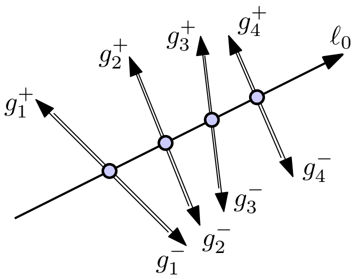

Figure 27: A minimal pinning of ℓ_0 by eight constraints.

The constant eight in Theorem 5.1 is best possible, as the following example shows. Four generically chosen non-oriented lines $g_1, \dots, g_4$ meeting ℓ_0 and perpendicular to it will have at most two common transversals: if no two among g_1, g_2, g_3 are coplanar or concurrent, their transversals define a hyperbolic paraboloid and it suffices to choose g_4 not lying on this surface. Now, let g_i^+ and g_i^- denote the two oriented lines supported by g_i (see Figure 27). Since a line satisfies g_i^+ and g_i^- if and only if it meets g_i, the eight constraints $g_1^+, g_1^-, \dots, g_4^+, g_4^-$ pin ℓ_0. Suppose we remove one of the eight constraints, say g_1^+. The common transversals to the three remaining lines g_2, g_3, g_4 form a quadric surface. By our construction, g_1 intersects this quadric transversely (since all four lines are orthogonal to ℓ_0, g_1 cannot be tangent to the quadric, and since the four lines have at most two transversals, g_1 cannot lie in the quadric). Thus, the quadric of transversals of g_2, g_3, g_4 contains lines on both sides of g_1, and ℓ_0 is no longer pinned. Therefore the eight oriented lines form a minimal pinning, as claimed.

⋆ ⋆ ⋆

We now introduce two parameterizations of the space of lines that are well-suited for analyzing the volume of lines satisfying one or several constraints.

We choose a coordinate system where ℓ_0 is the positive z-axis, and denote by $\mathcal{L}$ the family of lines whose direction vector makes a positive dot-product with $(0,0,1)$. Since pinning is a local property, we can decide whether ℓ_0 is pinned by considering only lines in $\mathcal{L}$. We identify $\mathcal{L}$ with $\mathbb{R}^4$ using the intersections of a line with the planes $z = 0$ and $z = 1$: the point $u = (u_1, u_2, u_3, u_4)$ in $\mathbb{R}^4$ represents the line $\ell(u)$ passing through the points $(u_1, u_2, 0)$ and $(u_3, u_4, 1)$ (see Figure 28). The line ℓ_0 is represented by the origin in $\mathbb{R}^4$, which we denote by O.

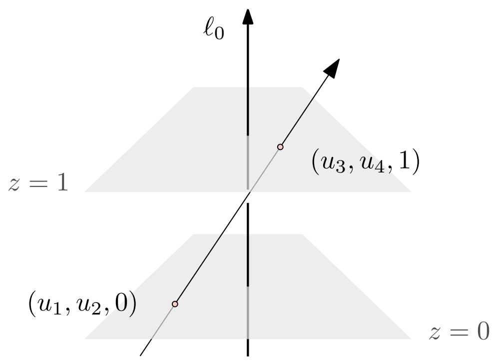

Figure 28: Parameterization of $\mathcal{L}$ by $\mathbb{R}^4$.

Our second model lifts $\mathcal{L}$ to a hypersurface of degree 2 in $\mathbb{R}^5$. Specifically, the map

$$\psi : \begin{cases} \mathbb{R}^4 & \rightarrow \quad \mathbb{R}^5 \\ (u_1, u_2, u_3, u_4) & \mapsto \quad (u_1, u_2, u_3, u_4, u_2u_3 - u_1u_4) \end{cases}$$

identifies $\mathcal{L}$ with the quadratic surface

$$\mathfrak{M} : \{u_5 = u_2 u_3 - u_1 u_4\}$$

in $\mathbb{R}^5$. Let $g(\lambda, \alpha, \delta)$ denote the constraint that meets ℓ_0 in the point $(0, 0, \lambda)$, makes slope δ with the plane $z = \lambda$, and projects into the xy-plane in a line making an angle α with the positive x-axis. Putting

$$\eta_g = \eta(\lambda, \alpha) = \begin{pmatrix} (1-\lambda)\sin\alpha \\ -(1-\lambda)\cos\alpha \\ \lambda\sin\alpha \\ -\lambda\cos\alpha \end{pmatrix}, \tag{5.2}$$

the volume of lines satisfying $g = g(\lambda, \alpha, \delta)$ writes, in the $\mathbb{R}^4$ parameterization,

$$\mathcal{U}_g = \{\delta(u_2 u_3 - u_1 u_4) + \eta_g \cdot u \leqslant 0\}.$$

If the constraint g is perpendicular to ℓ_0 then $\delta = 0$ and $\mathcal{U}_g$ is a halfspace with outer normal η_g; we call such constraints *orthogonal*. In general the solid $\mathcal{U}_g$ is bounded by a quadric through the origin. In our $\mathbb{R}^5$ parameterization, however, $\psi(\mathcal{U}_g) = \bar{\mathcal{U}}_g \cap \mathfrak{M}$ where

$$\bar{\mathcal{U}}_{g(\lambda,\alpha,\delta)} = \{\delta u_5 + \eta_g \cdot (u_1, \ldots, u_4) \leqslant 0\}$$

is a halfspace.

Consider, as a warm-up, the case where a family of orthogonal constraints pins ℓ_0. The volumes $\mathcal{U}_g$ are then halfspaces with O on their boundary, and a subfamily pins ℓ_0 if and only if the corresponding halfspaces intersect in exactly $\{O\}$. Now, a family of halfspaces each of which has O on its boundary intersects in exactly $\{O\}$ if and only if O lies in the interior of the convex hull of their outer normals. This suggests to use the following classic theorem of Steinitz:

Theorem 5.3 (Steinitz)**.** *If a point is interior to the convex hull of a set $X \subset \mathbb{R}^d$, then it is interior to the convex hull of some subset of at most $2d$ points of X.*

Steinitz's theorem implies that whenever a family of orthogonal constraints pin ℓ_0 some eight of them already pin ℓ_0. The bound $2d$ is only sharp if the convex hull of X has $2d$ vertices that form d pairs (x, x') with the surrounded point lying on the segment xx' [Rob42, Lemma 2a]. The example of Figure 27 is thus the *only* example of a minimal pinning by eight orthogonal constraints.

In the general case the set of lines in $\mathcal{L}$ satisfying all constraints in a family $\mathcal{F}$ is identified with

$$\bigcap_{g\in\mathcal{F}} (\bar{\mathfrak{U}}_g \cap \mathfrak{M}) = \left(\bigcap_{g\in\mathcal{F}} \bar{\mathfrak{U}}_g\right) \cap \mathfrak{M}, \tag{5.3}$$

which is the intersection of the quadratic hypersurface $\mathfrak{M}$ with a polyhedral cone C with apex the origin. Our goal is to show that there exists a subset of at most eight halfspaces in $\{\bar{\mathfrak{U}}_g \mid g \in \mathcal{F}\}$ that intersect in a cone C' such that the origin is an isolated point of $C' \cap \mathfrak{M}$.

⋆ ⋆ ⋆

Before we discuss the general case we make a few preliminary remarks on cones. A j-*space* denotes a j-dimensional linear subspace and a *cone* is the intersection of halfspaces whose bounding hyperplanes go through the origin. Unless specified otherwise all halfspaces are closed.

Let $\langle X\rangle$ denote the linear hull of a set $X \subset \mathbb{R}^d$, that is, the smallest linear subspace of $\mathbb{R}^d$ containing X. If H is a family of halfspaces in $\mathbb{R}^d$ and $K = \bigcap_H$ then

$$\langle K\rangle = \bigcap_{\substack{h\in H\\ \langle K\rangle\subseteq h}} h = \bigcap_{\substack{h\in H\\ \langle K\rangle\subseteq h}} \partial h, \tag{5.4}$$

and if K is contained in a halfspace h_0 we also have

$$\bigcap_{\substack{h\in H\\ \langle K\cap\partial h_0\rangle\subseteq h}} h \subseteq h_0. \tag{5.5}$$

(see [1, Lemma 5 and 7]).

We will also need the following Helly-type theorem for the condition of "intersecting a halfspace in exactly a point".

Lemma 5.4 ([1, Lemma 8]). *Let* $d \geqslant 2$ *and equip* $\mathbb{R}^d$ *with a coordinate system* $(O, x_1, \ldots, x_d)$. *If* H *is a family of closed halfspaces such that no* $h \in H$ *is bounded by* $\{x_d = 0\}$ *and* $\bigcap_H$ *is contained in* $\{O\} \cup \{x_d > 0\}$ *then there is a subfamily* $H' \subseteq H$ *of size at most* $2d-2$ *such that* $\bigcap_{H'}$ *is contained in* $\{O\} \cup \{x_d > 0\}$.

Proof. We first use Helly's theorem to find a subset $H_d \subseteq H$ of size at most d such that $C_d = \bigcap_{H_d}$ does not intersect $\{x_d = -1\}$ and is therefore contained in $\{x_d \geqslant 0\}$. Since no halfspace in H_d is bounded by $\{x_d = 0\}$, if $d = 2$ then $C_d \cap \{x_d = 0\}$ cannot be 1-dimensional and the statement holds. This establishes the induction basis for an inductive proof. Let $d > 2$ and assume that the statement holds for dimensions $2 \leqslant j < d$.

Let $E = \langle C_d \cap \{x_d = 0\}\rangle$ and note that, since no halfspace of H_d is bounded by $\{x_d = 0\}$, the dimension k of E satisfies $1 \leqslant k \leqslant d-2$. Since $(\bigcap_H) \cap E = \{O\}$, Steinitz's theorem (Theorem 5.3) implies that some subset $H_1 \subseteq H$ of size at most $2k$ satisfies $(\bigcap_{H_1}) \cap E = \{O\}$. Let $J = \{h \in H_d \mid E \subset h\}$; from Identities (5.4) (applied to $C_d \cap \{x_d = 0\}$) and (5.5) we have:

$$E = \bigcap\nolimits_J \cap \{x_d = 0\} \quad \text{and} \quad \bigcap\nolimits_J \subseteq \{x_d \geqslant 0\}.$$

Let π denote the orthogonal projection on $E^\perp$, the $(d-k)$-space orthogonal to E. Since $\pi(\bigcap_J)$ is contained in $\{O\} \cup \pi(\{x_d > 0\})$, the induction hypothesis implies that $\pi(\bigcap_{H_2})$ is contained in $\{O\} \cup \pi(\{x_d > 0\})$ for some subset $H_2 \subseteq J$ of cardinality at most $2d - 2k - 2$. This implies that $\bigcap_{H_2}$ is contained in $E \cup \{x_d > 0\}$ and $\bigcap_{H_1 \cup H_2}$ is contained in $\{O\} \cup \{x_d > 0\}$. Since $H_1 \cup H_2$ has cardinality at most $2d-2$, the inductive step is complete. □

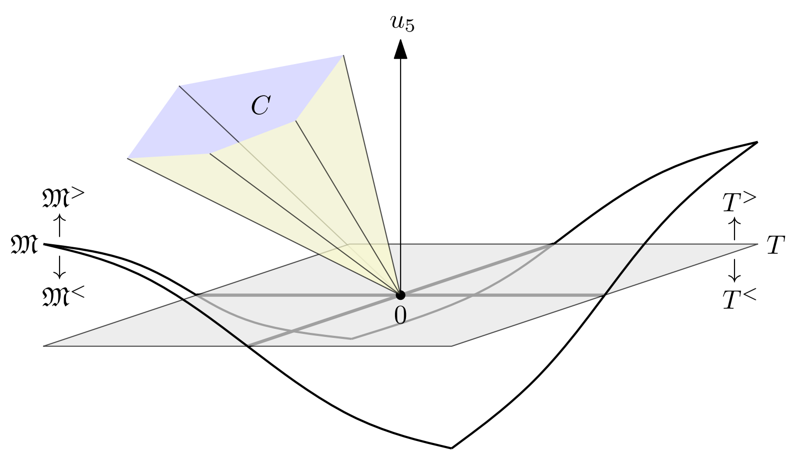

Figure 29: A symbolic drawing of the manifold $\mathfrak{M}$ and its tangent space T at o, and a cone C.

Our last ingredient is a characterization of isolated intersections between a cone and $\mathfrak{M}$. Let $T = \{u_5 = 0\}$ denote the hyperplane tangent to $\mathfrak{M}$ at the origin. Let $T^> = \{u_5 > 0\}$ and $T^< = \{u_5 < 0\}$ be the two open halfspaces bounded by T. Similarly, we denote the open regions above and below $\mathfrak{M}$ by $\mathfrak{M}^> = \{u_5 > u_2u_3 - u_1u_4\}$ and $\mathfrak{M}^< = \{u_5 < u_2u_3 - u_1u_4\}$ (see Figure 29). We define $T^\geqslant$, $T^\leqslant$, $\mathfrak{M}^\geqslant$, and $\mathfrak{M}^\leqslant$ analogously.

Lemma 5.5 ([1, Lemma 3]). *Let C be a cone in $\mathbb{R}^5$. The origin is an isolated point of $C \cap \mathfrak{M}$ if and only if either (i) C is a line intersecting T transversely, (ii) C is contained in $T^> \cup (T \cap \mathfrak{M}^>) \cup \{0\}$, or (iii) C is contained in $T^< \cup (T \cap \mathfrak{M}^<) \cup \{0\}$.*

Proof. Assume that the origin is isolated in $C \cap \mathfrak{M}$. We first observe that C intersects $T \cap \mathfrak{M}$ in exactly the origin as otherwise we could find a line through the origin in $C \cap T \cap \mathfrak{M}$. Then, assume, for the contradiction, that C contains both a point $u \in T^{>} \cup (T \cap \mathfrak{M}^{>})$ and a point $v \in T^{<} \cup (T \cap \mathfrak{M}^{<})$ and is not a line. We can ensure, by perturbing v if necessary, that the segment uv does not contain the origin. There is an $\varepsilon > 0$ such that for $t \in (0, \varepsilon)$, $tu \in \mathfrak{M}^{>}$ and $tv \in \mathfrak{M}^{<}$. The point w_t of $\mathfrak{M}$ on the segment joining tu and tv tends to the origin as t goes to 0, contradicting the assumption that the origin is isolated in $C \cap \mathfrak{M}$. The condition is therefore necessary.

A line intersecting a quadric transversely meets it in at most two points, so condition (i) is sufficient. Assume condition (ii) holds, as condition (iii) is symmetric. If $C = \{O\}$ we are done, otherwise $C \cap \mathfrak{M}^{\leqslant} \cap T = \{O\}$ and the sets

$$A = \{\, u \in C \mid \|u\| = 1 \,\} \quad \text{and} \quad B = \{\, u \in \mathfrak{M}^{\leqslant} \cap T \mid \|u\| \leqslant 1 \,\}$$

are disjoint. We let τ denote the distance between A and B; remark that $\tau > 0$ as both A and B are compact and non-empty. Now, given a point $u = (u_1, \dots, u_5) \in C \cap \mathfrak{M} \setminus \{0\}$ we define two points, $v = u/\|u\|$ and $v' = u'/\|u\|$ where $u' = (u_1, \dots, u_4, 0)$. Since $v \in A$ and $v' \in B$ (the latter following from $u_5 > 0$) the distance $\|v - v'\|$ is at least τ, implying that:

$$u_5 = \|u - u'\| = \|v - v'\| \cdot \|u\| \geqslant \tau \|u\|.$$

Since $u_5 = u_2 u_3 - u_1 u_4$ the classic inequality $xy \leqslant \frac{x^2+y^2}{2}$ implies that $\|u\|^2 \geqslant 2u_5$. Altogether, any point $u \in C \cap \mathfrak{M} \setminus \{0\}$ satisfies $\|u\| \geqslant 2\tau$, and the origin is isolated in $C \cap \mathfrak{M}$. □

⋆ ⋆ ⋆

We can now prove Theorem 5.1: we start with a family $\mathcal{F}$ of constraints that pins ℓ_0 and argue that some subset of at most eight of them already pins ℓ_0. We put $H = \{\bar{\mathfrak{U}}_g \mid g \in \mathcal{F}\}$ and let $C = \bigcap_H$. Since $\mathcal{F}$ pins ℓ_0, the origin is an isolated point of $C \cap \mathfrak{M}$ and we are in one of the cases (i)–(iii) of Lemma 5.5.

In case (i), C is a line. By considering the orthogonal projection of the $\bar{\mathfrak{U}}_g$ on the $\mathbb{R}^4$ perpendicular to C, we can use Steinitz's theorem (Theorem 5.3) to find a subset $H' \subseteq H$ of at most eight halfspaces such that $\bigcap_{H'} = C$. It follows that the at most eight constraints corresponding to H' suffice to pin ℓ_0.

Without loss of generality, we can now assume that we are in case (ii) of Lemma 5.5, that is $C \subseteq T^{>} \cup (T \cap \mathfrak{M}^{>}) \cup \{0\}$. Let

$E = \langle C \cap T \rangle$ and let k denote the dimension of E. We proceed in two steps: we first find a subset of at most $8-2k$ halfspaces in H that prevent C from entering $T^<$, then identify $2k$ additional halfspaces in H that, in the hyperplane T, prevent C from entering $\mathfrak{M}^{\leqslant}$. Lemma 5.5 (ii) then imply that the constraints corresponding to those at most 8 halfspaces already pin ℓ_0.

The first step uses the same idea as the proof of Lemma 5.4. Let π denote the orthogonal projection on $E^\perp$, the $(5-k)$-space orthogonal to E. Since T does not bound any halfspace in H, we have $E \neq T$ and so $0 \leqslant k \leqslant 3$. We put $H_1 = \{h \in H \mid E \subseteq h\}$. From Identities (5.4) (applied to $C \cap T$ inside the 4-space T) and (5.5) we have:

$$\bigcap_{H_1} \cap T = E \quad \text{and} \quad \bigcap_{H_1} \subseteq T^{\geqslant}.$$

Together this implies that $\pi\left(\bigcap_{H_1}\right) \subseteq \{O\} \cup \pi(T^>)$. Applying Lemma 5.4 in $E^\perp$, we have a subfamily $H_2 \subseteq H_1$ of size at most $2(5-k)-2 = 8-2k$ such that $\pi\left(\bigcap_{H_2}\right) \subseteq \{O\} \cup \pi(T^>)$, implying that $\bigcap_{H_2} \subseteq E \cup T^>$. This establishes the first step.

For the second step, we give a direct specific geometric argument in the k-dimensional subspace E, for each value of k. If $k = 0$ there is nothing to prove. If $k = 1$, then E is a line contained in T. If E intersects $\mathfrak{M}$ in a point other than the origin then, as $\mathfrak{M}$ is the graph of an homogeneous polynomial, $E \subset \mathfrak{M}$ and $\mathcal{F}$ cannot pin ℓ_0. It follows that $E \cap T = \{O\}$ and there is, again, nothing to prove.

If $k = 2$ then $C \cap T$ is a plane, a halfplane, or a convex wedge lying in the 2-space E. We can pick at most two constraints h_1, h_2 in H such that $h_1 \cap h_2 \cap T = C \cap T$. We then have that $\bigcap_{H_2 \cup \{h_1, h_2\}}$ is contained in $T^> \cup (T \cap C)$ which is contained in $T^> \cup (T \cap \mathfrak{M}^>) \cup \{0\}$, and we are done.

When $k = 3$ we first observe, via Identity 5.4, that $E = T \cap \partial\bar{\mathfrak{U}}_{g_0}$ with $g_0 \in \mathcal{F}$. Since $T \cap \mathfrak{M}$ corresponds to the lines intersecting ℓ_0, $E \cap \mathfrak{M}$ is the set of lines meeting ℓ_0 and g_0. We can then write $E \cap \mathfrak{M} = E_1 \cup E_2$ where E_1 is the set of lines through $\ell_0 \cap g_0$ and E_2 the set of lines in the plane spanned by ℓ_0 and g_0. Both E_i turn out to be 2-dimensional planes, and together they partition the 3-space E into four quadrants. Since $C \cap T$ intersects $\mathfrak{M}$ only in the origin, it must be contained in one of these quadrants. The projection of $C \cap T$ along $E_1 \cap E_2$ is a two-dimensional wedge bounded by the projections of two edges of the three-dimensional cone $C \cap T$. Each edge is defined by at most two constraints of $C \cap T$, and thus we can find at most four constraints H_3 of $C \cap T$ that define the same projected wedge. This ensures that

$\bigcap_{H_2 \cup H_3}$ is contained in $T^> \cup (T \cap \mathfrak{M}^>) \cup \{0\}$, and we are done. This concludes the proof of Theorem 5.1.

⋆ ⋆ ⋆

We finally come to the construction of arbitrarily large minimal pinnings of a line by polytopes in $\mathbb{R}^3$, proving Theorem 5.2.

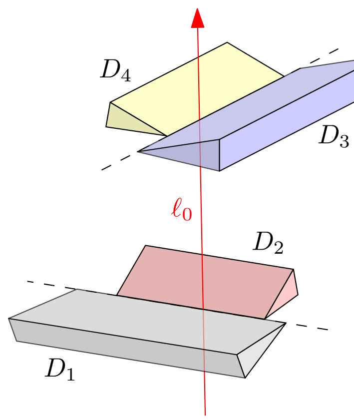

Figure 30: Four polytopes with a two-dimensional set of transversals.

We first pick two polytopes D_1 and D_2 such that their common transversals in the vicinity of ℓ_0 are precisely the lines intersecting the y-axis. Similarly, we pick two polytopes D_3 and D_4 that restrict the transversals to pass through the line $\{(t, 0, 1) \mid t \in \mathbb{R}\}$, as in Figure 30. A line $\ell(u)$ meets all four polytopes if and only if $u_1 = u_4 = 0$. We can therefore analyze the situation in the u_2u_3-plane. We add two other polytopes D_5 and D_6 (not pictured) to enforce $2u_2 \geqslant u_3$ and $2u_3 \geqslant u_2$; these polytopes are bounded by the oriented lines $\{(t, -t, -1) \mid t \in \mathbb{R}\}$ and $\{(t, -t, 2) \mid t \in \mathbb{R}\}$. In the u_2u_3-plane, the set of lines meeting $D_1, \ldots, D_6$ is the closed wedge $W = \{(u_2, u_3) \mid u_2/2 \leqslant u_3 \leqslant 2u_2\}$.

Consider two angles, β and θ, with $0 < \beta < \theta < \pi/2$. Let $v = (v_x, v_y) = (\cos\beta, \sin\beta)$, $w = (w_x, w_y) = (\cos\theta, \sin\theta)$ be two unit vectors, and define the unbounded polyhedral wedge:

$$F(v, w) = \{(x, y, z) \mid v_x x + v_y y \leqslant 0 \text{ and } w_x x + w_y y \leqslant 0\}.$$

The left-hand side of Figure 31 shows a projection along ℓ_0. A line $\ell(u)$ with $u_2, u_3 > 0$ and $u_1 = u_4 = 0$ misses $F(v, w)$ if and only if the vector $(u_2, u_3) \in \mathbb{R}^2$ falls in the (closed, counterclockwise) acute angular interval $\xi_{v,w} = [\beta, \theta]$ between v and w. In other words, the set of lines intersecting $F(v, w)$ is the u_2u_3-plane with the closed wedge corresponding to $\xi_{v,w}$ removed (see the blue shape in Figure 31 (right)).

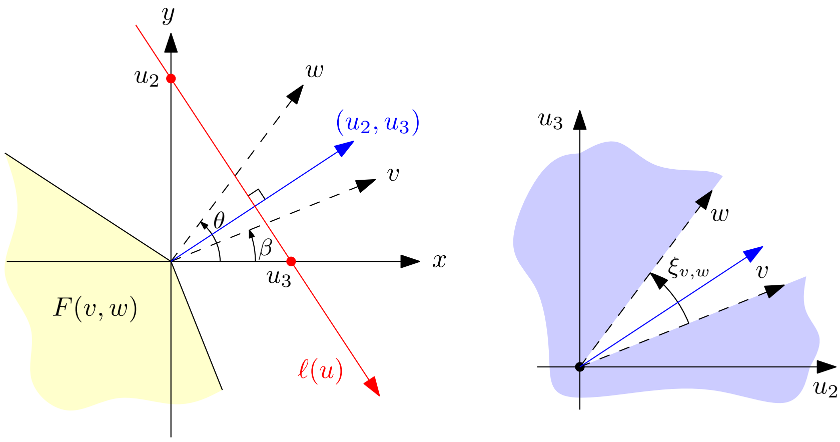

Figure 31: The polyhedron $F(v, w)$ and the set of lines intersecting it.

Let $(v^1, w^1), \ldots, (v^n, w^n)$ be n pairs of vectors such that together the wedges ξ_{v^i,w^i} cover W and each middle vector $v^i + w^i$ of $\xi_{v^i w^i}$ lies in W but not in any $\xi_{v^j w^j}$ with $j \neq i$. Now, ℓ_0 is the only line in W intersecting every $F(v^i, w^i)$ but for any $1 \leqslant i \leqslant n$ there is an entire sector of lines in W that intersect all F_j with $j \neq i$. It follows that the family $\mathcal{F} = \{D_1, \ldots, D_6, F(v^1, w^1), \ldots, F(v^n, w^n)\}$ pins ℓ_0, but has no pinning subfamily of size smaller than n (some D_i could be redundant, but none of the $F(v^i, w^i)$ is).

Since pinning is determined by lines in a neighborhood of ℓ_0 only, we can clearly make $D_1, \ldots, D_6$ bounded and it only remains to crop the $F(v^i, w^i)$. For $t \in [0, 1]$ we let $\gamma_i(t)$ be the line whose parameters in $\mathbb{R}^4$ are $(0, v^i_y + w^i_y, v^i_x + w^i_x, 0) \cdot t$. Since $\gamma_i(t)$ starts at the origin and moves, in the $(u2, u3)$-plane, on a line perpendicular to the vector $v^i + w^i$, it misses $F(v^i, w^i)$ while intersecting each $F(v^j, w^j)$ with $j \neq i$. A straightforward calculation shows that the point where $\gamma_i(t)$ enters or exits F_j moves linearly away from ℓ_0 along a line perpendicular to ℓ_0. We can therefore crop each F_j to a bounded polytope, ensuring that it still contains all points $P_{ij}(t)$ for $i \neq j$ and $0 < t \leqslant 1$, and hence intersects all lines on the paths $\gamma_i(t)$ for $i \neq j$. As a result, the family $\mathcal{F}$ still pins ℓ_0, but for any F_i the family $\mathcal{F} \setminus \{F_i\}$ is not a pinning as witnessed by the path γ_i. This completes the proof of Theorem 5.2.

⋆ ⋆ ⋆

NOTES.

This chapter is based on an article co-authored with Boris Aronov, Otfried Cheong and Günter Rote [1].

Our manifold $\mathfrak{M}$ is the image of the Klein quadric under a mapping that sends (the point of the Klein quadric representing) ℓ_0 to the origin, the hyperplane tangent to the Klein quadric at ℓ_0 to $\{u_5 = 0\}$, and the lines orthogonal to ℓ_0 to infinity. Our affine representation has the advantage that $\mathfrak{M}$ admits a parameterization of the form $u_5 = f(u_1, \ldots, u_4)$ where f is a homogeneous polynomial of degree two. This is instrumental in the proof of Lemma 5.5. Various other properties we used are well-known in the Klein model.

Holmsen and Matoušek [HM04] constructed a family of examples showing that the transversal Helly number of families of disjoint translates of a convex polytope in $\mathbb{R}^3$ cannot be bounded. Thus, while the absence of a line transversal *globally* cannot be witnessed by a small certificate, Theorem 5.1 asserts that locally this is the case.

As mentioned previously (page 34), the notion of pinning is a natural counterpart to the notion of grasping, or immobilizing, studied in robotics. We give two examples of how Theorem 5.1 naturally translates into Helly numbers for grasping. First, we can interpret a family $\mathcal{F}$ of lines pinning a line ℓ_0 by considering all lines as solid cylinders of zero radius. Each constraint "cylinder" touches the "cylinder" ℓ_0 on the left. As a result, it is impossible to move ℓ_0 in any way (except to rotate it around or translate it along its own axis), because it would then intersect one of the constraints, so ℓ_0 is grasped by $\mathcal{F}$. Second, consider a family $\mathcal{F} = \{P_1, \ldots, P_n\}$ of polytopes such that P_i is tangent to ℓ_0 in a single point interior to an edge e_i, and let $\tilde{P_i}$ denote the mirror image of P_i with respect to the plane spanned by ℓ_0 and e_i. Since $\mathcal{F}$ pins ℓ_0 if and only if $\tilde{\mathcal{F}}$ grasps it, our pinning theorems directly translate into local Helly numbers for grasping a line by polytopes.

6 PINNING BY CONSTRAINTS

In this chapter, we supplement the pinning theorem of Chapter 5 by an exploration of the various configurations that constitute a minimal pinning of a line by constraints in $\mathbb{R}^3$.

We, again, fix a line ℓ_0 to be pinned and use the parameterization of lines not orthogonal to ℓ_0 by $\mathbb{R}^4$ introduced on page 53. In this setup, the solid of lines satisfying an orthogonal constraint is a halfspace with the origin on its boundary. Such a family intersects in a single point if and only if the convex hull of their outer normal vectors contains the origin in its interior. We give a description of minimal families of vectors (or, equivalently, points) in $\mathbb{R}^4$ containing the origin in the interior of their convex hull. We then characterize which of these configurations can be realized by normals of orthogonal constraints; a tabulation of all 16 types of minimal pinnings of a line by *orthogonal* constraints follows.

We then turn our attention to general constraints. We first discuss the cases where not only the constraints pin ℓ_0, but the first-order approximations of the solids $\mathcal{U}_g$ also intersect in a single point; these "first-order pinnings" are naturally related to pinnings by orthogonal constraints. We conclude this chapter by a few considerations on "higher-order pinnings".

⋆ ⋆ ⋆

We say that a set of points in $\mathbb{R}^d$ *surrounds the origin* if the origin lies in the interior of their convex hull. A set S *minimally surrounds* the origin if S surrounds the origin, but no proper subset of S does. Our first goal is to describe sets minimally surrounding the origin in $\mathbb{R}^4$ as unions of (not necessarily disjoint) critical simplices. A *simplex* of dimension k, or k-*simplex*, is a set of $k+1$ affinely independent points in $\mathbb{R}^d$ (we also say *segment*, *triangle* and *tetrahedron* for $k = 1, 2$ and 3). We call a simplex *critical* if it surrounds the origin in its linear hull. Equivalently, σ is critical if and only if every point $y \in \langle\sigma\rangle$ can be written as $y = \sum_{x\in\sigma} \lambda_x x$ with all $\lambda_x \geqslant 0$; we also say that σ *positively spans* $\langle\sigma\rangle$.

We prove [1, Theorem 5] that a set S minimally surrounds the origin in $\mathbb{R}^4$ if and only if the linear hull of S is $\mathbb{R}^4$ and one of the following holds:

(i) $|S| = 5$ and S is a critical 4-simplex, or

(ii) $|S| = 6$ and S is the union of two critical simplices, each of dimension at most three, or

(iii) $|S| = 7$ and S is the union of three critical simplices: $k \geqslant 1$ critical triangles having a single point in common and $3 - k$ disjoint critical segments, or

(iv) $|S| = 8$ and S is the disjoint union of four critical segments.

The cases (ii)-(iv) are represented in Figure 32. Note that we are not claiming that the critical simplices shown are *all* critical simplices of the point set (although we are not aware of a situation that has additional critical simplices).

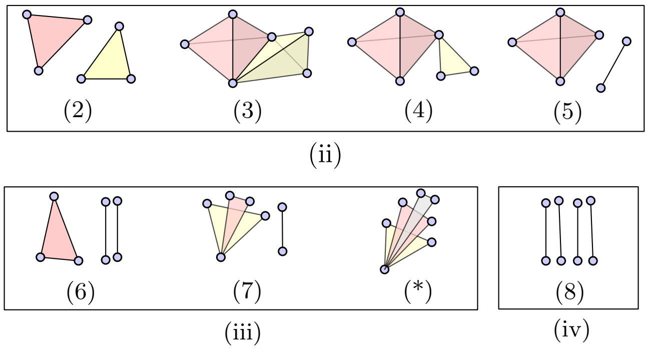

Figure 32: A combinatorial description of nongeneric minimal sets of points surrounding the origin in $\mathbb{R}^4$.

Let us first argue that the cases (i) to (iv) exhaust all possibilities. Let S be a minimal set of points that surrounds 0. We first remark that Caratheodory's theorem guarantees that S contain some critical simplex [1, Lemma 10]. We pick A to be a critical simplex of maximum dimension contained in S. If $|A| = 5$ then we are in case (i) and if $|A| = 2$ any critical simplex of S has size exactly two and it is easy to see that we must be in case (iv). Otherwise we let $B = S \setminus A$ and let π denote the orthogonal projection on $\langle A \rangle^{\perp}$; we denote by $conv(X)$ the convex hull of X. We use three elementary properties [1, Lemma 12]:

(a) $A \cup B$ surrounds the origin if and only if $\pi(B)$ surrounds the origin in $\langle A \rangle^{\perp}$.

(b) if $A \cup B$ *minimally* surrounds the origin then $\pi(B)$ *minimally* surrounds the origin in $\langle A \rangle^{\perp}$.

(c) If $conv(B) \cap \langle A \rangle \neq \emptyset$, and $conv(X) \cap \langle A \rangle = \emptyset$ for every $X \subsetneq B$, then B is contained in a critical simplex of $A \cup B$.

If $|A| = 4$ then $\langle A \rangle^{\perp}$ is a line and (b) implies that B consists of two points, one on each side of $\langle A \rangle$; property (c) then implies that we are in case (ii).

Now, if $|A| = 3$ then $\langle A \rangle^{\perp}$ is a 2-plane and $\pi(B)$ consists of a critical triangle or two critical segments. If $\pi(B)$ is a critical triangle then the affine hull of B intersects $\langle A \rangle$ in a single point interior to the convex hull of B; also, no edge of the triangle B meets $\langle A \rangle$, so (c) ensures that B is contained in a critical simplex of S. That simplex has cardinality at most $|A| = 3$, so B is a critical triangle and we are in case (ii). Assume now that $\pi(B)$ is two critical segments and let B_1 and B_2 be the two corresponding segments in B. Again, property (c) ensures that each B_i is contained in a critical simplex of S. Since A is of maximal cardinality, B_i is a critical segment or is contained in a critical triangle $T_i = B_i \cup \{a_i\}$. If at least one B_i is a critical segment, then we are in case (iii) (case (6) or (7) of Figure 32). If both B_i are contained in a critical triangle and $a_1 \neq a_2$ then $\langle\{a_1, a_2\}\rangle = \langle A \rangle$ and $\langle T_1 \cup T_2 \rangle = \mathbb{R}^4$; (a) then implies that $T_1 \cup T_2$ surrounds the origin in $\mathbb{R}^4$, contradicting the minimality of S. We are thus in case (iii) (case (*) of Figure 32). This proves that any point set minimally surrounding the origin in $\mathbb{R}^4$ is of one of the types (i)–(iv).

What about the converse? Assume that S is a union of critical simplices and that $\langle S \rangle = \mathbb{R}^4$. Let h be any closed halfspace containing the origin on its boundary. Since $\langle S \rangle = \mathbb{R}^4$ some point p of S lies in the interior of h or outside h. In the former case, a critical simplex of S that contains p cannot be contained in h, and therefore also contains a point outside h. It follows that any closed halfspace containing the origin on its boundary misses at least a point of S. It follows that S surrounds the origin. Arguing that each of the configurations (i)–(iv) *minimally* surrounds is more tedious and requires separate arguments for each case. We omit the details here and refer, instead, to [1, Therorem 5].

⋆ ⋆ ⋆

Our next step is to characterize geometrically the situations where a family $\{\eta_{g_0}, \ldots, \eta_{g_k}\}$ of normals to orthogonal constraints form a critical simplex.[1]

There are two necessary conditions for a family $\mathcal{F} = \{u_0, \ldots, u_k\}$ of vectors in $\mathbb{R}^4$ to form a critical simplex. First, $\mathcal{F}$ must be linearly dependent (so that the origin lies in its linear hull) and any proper subset of $\mathcal{F}$ must be linearly independent. Second,

1 We say that a family of vectors $u_0, \ldots, u_k$ forms a critical simplex if the points $O + u_0, \ldots, O + u_k$ form a critical simplex.

the intersection of the halfspaces with the origin on their boundary and outer normals $u_0, \dots, u_k$ should have dimension exactly $4-k$. These conditions are easily checked to be sufficient. If the u_i are normal vectors to constraints $\mathcal{U}_{g_i}$ then the two conditions are equivalent to:

(i) the dimension of the space of lines satisfying $g_0, \dots, g_k$ must be $4-k$, and

(ii) every constraint must meet all the lines meeting all the other constraints.

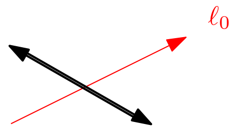

Figure 33: A 2-block, giving rise to a critical segment.

Consider first two orthogonal constraints g_0 and g_1. Since every line meeting g_1 also meets g_2, the constraints are equal or are the two orientations of the same unoriented line. The former case leads to $\eta_{g_0} = \eta_{g_1}$, which is not a critical segment, whereas the latter case leads to $\eta_{g_0} = -\eta_{g_1}$, which is a critical segment. We call the two opposite orientations of a line orthogonal to ℓ_0 a *2-block*.

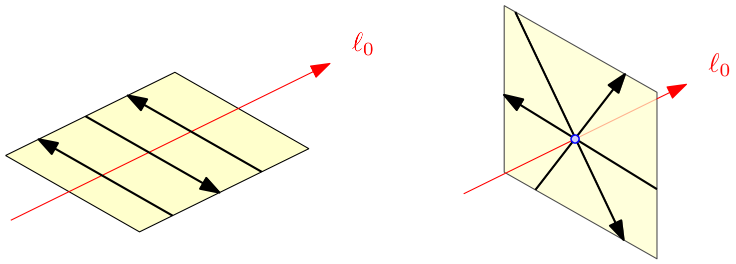

Figure 34: The 3-blocks giving rise to critical triangles: $3^{\parallel}$-block (left) and $3^{\times}$-block (right).

Consider now three orthogonal constraints g_0, g_1, g_2. The condition that every line meeting two constraints also meets the third forbids any two of the constraints to be skew. Since the g_i are orthogonal constraints, we are in one of two cases:

- g_0 and g_1 can be coplanar with ℓ_0. In that case, g_2 has to lie in the same plane as well. The set E of lines satisfying the three constraints must be two-dimensional; E already contains the set of lines in the plane spanned by the g_i, which is two-dimensional. This imposes that any two g_i met consecutively by ℓ_0 have alternating orientations; we call such a triple of constraints a $3^{\parallel}$-*block* (see Figure 34-left).

- g_0 and g_1 can meet in a point $p \in \ell_0$. In that case, g_2 goes through p as well. The set E of lines satisfying the three constraints must be two-dimensional; E already contains the set of lines through p, which is two-dimensional. This imposes that the direction vectors of the g_i positively span $\ell_0^\perp$; we call such a triple of constraints a $3^\times$-*block* (see Figure 34-right).

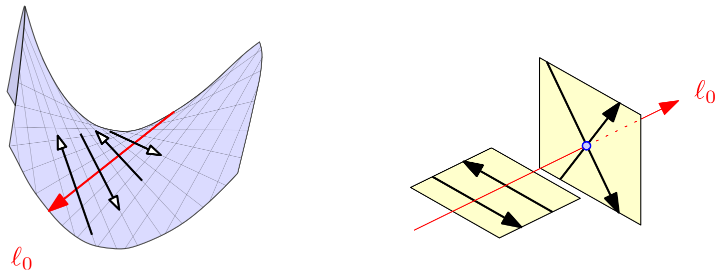

Figure 35: The 4-blocks, giving rise to critical tetrahedra: $4^\parallel$-block (left) and $4^\times$-block (right).

Consider four orthogonal constraints g_0, g_1, g_2, g_3. Any three constraints have a one-dimensional family of common transversals, so we are in one of two cases:

- the g_i are pairwise skew. Then g_3 must lie in the hyperbolic paraboloid formed by the transversals to g_0, g_1, g_2; we also say that the lines are in *hyperboloidal position*. The set E of lines satisfying the four constraints must be one-dimensional; E already contains the lines in the other family of rulings of the quadric containing the g_i. It follows that the constraints must be oriented such that only lines lying in the quadric satisfy all four constraints; we call such a quadruple of constraints a $4^\parallel$-*block* (see Figure 35-left).

- two of the constraints are coplanar or concurrent with ℓ_0; say g_0 and g_1 meet in $p \in \ell_0$. The remaining constraints cannot contain p as otherwise three constraints have a two-dimensional set of line transversals. The condition that g_3 meets any line meeting g_0, g_1, g_2 implies that g_2 and g_3 are coplanar with ℓ_0; conversely, if we assume that g_0 and g_1 are coplanar with ℓ_0 similar arguments yield that g_2 and g_3 need to be concurrent with ℓ_0. Let p be the intersection point of the concurrent pair and Π be the plane spanned by the coplanar pair. The set E of lines satisfying the four constraints must be one-dimensional; E already contains the lines through p and contained in Π, which is one-dimensional. It follows that the constraints must be oriented such that only lines through p in Π satisfy all four constraints; we call such a quadruple of constraints *a $4^{\times}$-block* (see Figure 35-right).

Finally, the normals of five constraints $g_0, \dots, g_4$ form a critical 4-simplex if and only if ℓ_0 is the only line satisfying all of them. We call such a family of constraints a 5-*block*. Interestingly, given five orthogonal constraints such that no four are dependent, we can always orient them (that is, reverse some of them) so as to obtain a 5-block.

Combining this characterization of critical simplices of normals with our previous description of minimal point sets surrounding the origin, we obtain the characterization of minimal pinnings of a line by orthogonal constraints summarized in Figure 38 (page 72). That each configuration pins ℓ_0 is straightforward. The minimality can be established by simple, if pedestrian, arguments specific to each configuration [1, Theorem 6].

⋆ ⋆ ⋆

We now turn our attention to situations where constraints that are not necessarily orthogonal pin "at first order", that is, when the origin remains isolated in the intersection of *the linearizations* of the solids $\mathcal{U}_g$.

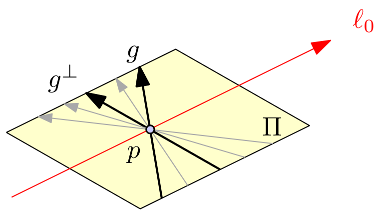

Figure 36: A constraint g and its orthogonalization $g^{\perp}$. All constraints represented have the same normal.

Let g be a constraint. The boundary of $\mathcal{U}_g$ is a quadric through the origin and its normal η_g in the origin is given by Equation 5.2. Let p denote the point $g \cap \ell_0$ and Π the plane spanned by g and ℓ_0. Rotating g around p inside Π, by varying the parameter δ, doesn't change the normal η_g. The linearization of the volume $\mathcal{U}_g$ is therefore the solid $\mathcal{U}_{g^{\perp}}$, where $g^{\perp}$ denotes the projection of g on the plane perpendicular to ℓ_0 in $\ell_0 \cap g$; we call $g^{\perp}$ the *orthogonalized constraint* of g, and denote by $\mathcal{F}^{\perp}$ the family of orthogonalized constraints of $\mathcal{F}$. Note that $\mathcal{F}^{\perp}$ can have smaller cardinality than $\mathcal{F}$.

Lemma 6.1. *Let $\mathcal{F}$ be a family of constraints. If $\mathcal{F}^{\perp}$ pins ℓ_0 then $\mathcal{F}$ pins ℓ_0. If $\mathcal{F}$ pins ℓ_0 and no four constraints in $\mathcal{F}$ have linearly dependent normals then $\mathcal{F}^{\perp}$ pins ℓ_0.*

Proof. Let $\mathcal{F}$ be a family of constraints. Since the sets $\mathcal{U}_g$ are bounded by algebraic surfaces of constant degree, the origin 0 is isolated in the intersection of such volumes if and only if there exists no smooth path moving away from 0 inside that intersection. Moreover, if the tangent vector at 0 to a smooth path γ makes a positive dot product with η_g, then γ locally exits $\mathcal{U}_g$. Now, if $\mathcal{F}^\perp$ pins ℓ_0 then, $\{\eta_g \mid g \in \mathcal{F}^\perp\} = \{\eta_g \mid g \in \mathcal{F}\}$ surrounds the origin and any vector must make a positive dot product with the normal to at least one of the constraints in $\mathcal{F}$, and $\mathcal{F}$ also pins ℓ_0. The same argument shows that if $\mathcal{F}$ pins ℓ_0, then $\bigcap_{g \in \mathcal{F}} \mathcal{U}_{g^\perp}$ must have empty interior. In that case, if it is not a single point then four of the normals to the constraints in $\mathcal{F}^\perp$ are linearly dependent. □

If $\mathcal{F}^\perp$ pins ℓ_0 we say that $\mathcal{F}$ pins ℓ_0 *at first order*; if $\mathcal{F}$ pins ℓ_0 but $\mathcal{F}^\perp$ does not we speak of *higher-order pinning*. All minimal first-order pinnings can be obtained from Figure 38 by rotating each constraint so as to change its δ-parameter arbitrarily.

⋆ ⋆ ⋆

Identifying $\mathcal{L} \simeq \mathbb{R}^4$ with the hyperplane $T : \{u_5 = 0\}$ reveals an interesting connection between our two parameterizations: for any constraint g, the set $\bar{\mathcal{U}}_g$ intersects $\mathfrak{M}$ in the set of lines satisfying g and intersects T in the set of lines satisfying $g^\perp$. More formally, putting $\phi : (u_1, \dots u_4) \mapsto (u_1, \dots u_4, 0)$ we have:

$$\bar{\mathcal{U}}_g \cap T = \phi(\mathcal{U}_{g^\perp}). \tag{6.1}$$

This allows to sharpen our pinning theorem for the case of higher-order pinnings.

Theorem 6.2. *Any higher-order minimal pinning of a line by constraints in $\mathbb{R}^3$ has size at most six.*

Proof. We proved our pinning theorem for general constraints by analyzing the intersection of the cone $C = \bigcap_{g \in \mathcal{F}} \bar{\mathcal{U}}_g$ with $\mathfrak{M}$ near the origin in terms of the trace of C on the hyperplane T. This trace can be identified, via Equation (6.1), with the set of lines satisfying $\mathcal{F}^\perp$. The higher-order pinnings thus correspond to the cases $k = 1, 2, 3$ (page 58) where the minimal pinning alway has size at most 6. □

A family $\mathcal{F}^\perp$ has linearly dependent normals if some of its constraints form, up to reversing orientation, a 2-, 3- or 4-block (see Figure 33, 34 and 35). This characterization can be pulled back to $\mathcal{F}$ by observing that (i) two constraints are coplanar (resp. concurrent) with ℓ_0 if and only if their orthogonalized constraints are coplanar resp. concurrent) with ℓ_0, and (ii) $\mathcal{F}^\perp$ is in hyperboloidal position if and only if the constraints in $\mathcal{F}$ are pairwise skew and one constraint is tangent to the quadratic surface swept by line transversals to the other three constraints. Altogether we obtain that:

- Two constraints have linearly dependent normals if and only if they are at the same time coplanar and concurrent with ℓ_0.
- Three constraints have linearly dependent normals if and only if they are coplanar or concurrent with ℓ_0.
- Four constraints have linearly dependent normals if and only if (a) two are concurrent with ℓ_0 and the other two are coplanar with ℓ_0, or (b) one is tangent to the quadric formed by the transversals to the three others.

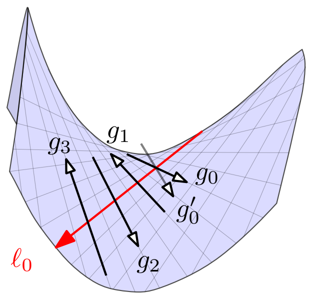

Figure 37: A (higher-order) minimal pinning of a line by 4 constraints.

We can now give a simple example of pinning of higher order. Start with a $4^{\parallel}$-block, for instance the constraints

$$\begin{aligned} g_a &= \{\,(-t,0,0) \mid t \in \mathbb{R}\,\}, \\ g_b &= \{\,(t,t,1) \mid t \in \mathbb{R}\,\}, \\ g_c &= \{\,(-t,-2t,2) \mid t \in \mathbb{R}\,\}, \text{ and} \\ g_d &= \{\,(t,3t,3) \mid t \in \mathbb{R}\,\}, \end{aligned}$$

oriented in the direction of increasing t. These constraints lie on the quadric $(\mathcal{B}) : y = xz$ and are simultaneously satisfied only by the rulings of $\mathcal{B}$ from the other family. Now rotate g_a into the line $g_a' = \{(-t,0,-t/100)\}$ tangent to $\mathcal{B}$ in the origin and otherwise contained in the volume $y < xz$ (see Figure 37). In order to satisfy g_a', g_b, g_c, g_d, a line near ℓ_0 would have to intersect $\mathcal{B}$ at least three times (points of tangency counted twice), and thus lies in $\mathcal{B}$ as $\mathcal{B}$ is a quadric; any line contained in $\mathcal{B}$ other than ℓ_0 violates g_0', so $\mathcal{F} = \{g_a', g_b, g_c, g_d\}$ pins ℓ_0. Its orthogonalization $\mathcal{F}^\perp = \{g_a, g_b, g_c, g_d\}$ does not pin ℓ_0.

It can be shown if $\mathcal{F}$ is a higher-order pinning of ℓ_0 then either two constraints in $\mathcal{F}$ are concurrent or coplanar with ℓ_0, or $\mathcal{F}^\perp$ is a $4^{\parallel}$-block [1, Theorem 7]. Any "non-degenerate" higher-order pinning is thus similar to the example above.

⋆ ⋆ ⋆

NOTES.

This chapter is based on an article co-authored with Boris Aronov, Otfried Cheong and Günter Rote [1].

As discussed pages 34 and 61, a family of lines pinning a line can be considered as a grasp of that line. In grasping, one often considers *form closure*, which means that the object is immobilized even with respect to infinitesimally small movements. For instance, an equilateral triangle with a point finger at the midpoint of every edge is immobilized, as it cannot be moved in any way, but it is not in form closure because an infinitesimal rotation around its center is possible. It is easy to see that all grasps listed in Table 38 are form closure grasps in this sense. The grasp caused by our example of higher-order pinning, however, is not a form-closure grasp, as ℓ_0 can be moved infinitesimally in the quadric defined by three of the lines.

In three dimensions, the set of lines in $\mathcal{L}$ intersecting a screen (B, ℓ_0) (as defined page 31) is precisely the set of lines satisfying the orthogonal constraint tangent to B in $B \cap \ell_0$. The three-dimensional case of Lemma 3.9, on the first-order approximation of transversals to a ball by transversals to a screen, is thus very similar to the statement that $\mathcal{U}_{g^\perp}$ is the first-order approximation of $\mathcal{U}_g$.

In the proof of our pinning theorem for polytopes (Chapter 5, page 58) we analyzed the case where the cone $\bigcap_{g \in \mathcal{F}} \bar{\mathcal{U}}_g$ intersects the plane $u_5 = 0$ in a k-dimensional face, for $k = 0, 1, 2, 3$. Using the interpretation of that face as the set of line satisfying $\mathcal{F}^\perp$, it is not difficult to construct minimal (higher-order) pinnings $\mathcal{F}$ realizing each of the cases $k = 1, \dots, 3$ [1, Section 5].

(1) A single 5-block;

(2a) Two disjoint $3^{\parallel}$-blocks defining distinct planes;

(2b) Two disjoint $3^{\times}$-blocks meeting ℓ_0 in distinct points;

(3a) Two $4^{\parallel}$-blocks sharing two constraints and defining distinct quadrics;

(3b) Two $4^{\times}$-blocks sharing two constraints, such that their coplanar pairs define distinct planes or their concurrent pairs define distinct points;

(3c) A $4^{\parallel}$-block and a $4^{\times}$-block sharing two constraints;

(4a) A $4^{\parallel}$-block and a $3^{\parallel}$-block sharing one constraint;

(4b) A $4^{\parallel}$-block and a $3^{\times}$-block sharing one constraint;

(4c) A $4^{\times}$-block and a $3^{\parallel}$-block sharing one constraint such that they define distinct planes;

(4d) A $4^{\times}$-block and a $3^{\times}$-block sharing one constraint such that their concurrent pairs meet ℓ_0 in distinct points;

(5a) A $4^{\parallel}$-block and a disjoint 2-block, where the 2-block constraints are not contained in the quadric defined by the $4^{\parallel}$-block;

(5b) A $4^{\times}$-block and a disjoint 2-block, where the 2-block constraints are neither coplanar with the coplanar pair nor concurrent with the concurrent pair of the $4^{\times}$-block;

(6a) A $3^{\parallel}$-block and two 2-blocks, where the four 2-block constraints do not all meet, and where no 2-block constraint is contained in the plane defined by the $3^{\parallel}$-block;

(6b) A $3^{\times}$-block and two 2-blocks, where the four 2-block constraints do not all meet, and where no 2-block constraint goes through the common point of the $3^{\times}$-block;

(7) A $3^{\parallel}$-block and a $3^{\times}$-block sharing one constraint, and a disjoint 2-block that does not lie in the plane of the $3^{\parallel}$-block and does not go through the common point of the $3^{\times}$-point;

(8) Four disjoint 2-blocks whose supporting lines are not in hyperboloidal position (that is, all orientations of four lines with finitely many common transversals).

Figure 38: The 16 types of minimal pinning of a line by orthogonal constraints in $\mathbb{R}^3$. The numbering of cases corresponds to the cases in Figure 32.

7 PINNING BY SMOOTH CONVEX SETS

In this chapter, we show that the previous analysis of pinnings by constraints easily extends to a full classification of minimal stable pinnings by smooth convex sets. We say that $\mathcal{F} = \{A_1, \ldots, A_n\}$ is a *stable pinning* of a line ℓ_0 if $\mathcal{F}$ remains a pinning after the A_i have been subject to sufficiently small (independent) screws[1] of axis ℓ_0.

We consider the solids $\mathcal{T}(A_i)$ from a differential point of view in an adequate representation of the space of lines. We show that when the ∂A_i are ovaloids, i.e. smooth convex surfaces with non-vanishing Gaussian curvature, the set $\partial\mathcal{T}(A_i)$ of lines tangent to A_i is a smooth manifold. We further interpret the tangent space to $\partial\mathcal{T}(A_i)$ in ℓ_0 as the set of lines intersecting some constraint[2] orthogonal to ℓ_0. This reduces the problem of classifying "minimal first-order pinnings" to the tabulation of Figure 38. We conclude by observing that for convex sets bounded by ovaloids, these "first-order pinnings" are exactly the stable pinnings, an idea already apparent in the proof of Theorem 3.10.

We also show, using arguments more in line with Chapters 2 and 3, a pinning theorem for families of convex sets bounded by ovaloids and satisfying a mild general position assumption: that no two sets be externally tangent in a point of the pinned line. Contrary to the pinning theorem obtained in Chapter 3, all our arguments in this chapter are *local*: we only need the sets to be convex, smooth, and have non-vanishing Gauss curvature in the vicinity of its contact point with the pinned line.

⋆ ⋆ ⋆

Before we discuss the smoothness of $\partial\mathcal{T}(\cdot)$ a few definitions are in order. Let A be a compact convex set in $\mathbb{R}^3$. A *support plane* of A is a plane that intersects A and bounds a closed halfspace containing A. A line *tangent to* A is a line that intersects A and is contained in one of its support planes. Let p be a point in ∂A. We say that A is *of class* C^k *in* p if there exists a neighborhood U of p in $\mathbb{R}^3$ such that $U \cap \partial A$ is a C^k-manifold. We say that A is

1 A *screw* of axis ℓ is the composition of a translation parallel to ℓ and a rotation of axis ℓ. By "sufficiently small" screw we mean screws where the angle of the rotation and the norm of the translation vector are sufficiently small.

2 A constraint here, as in Chapters 5 and 6, is simply an oriented line meeting ℓ_0.

rotund in p if A is C^2 in p and the Gaussian curvature of ∂A in p is positive. If A is rotund in p then ∂A is *strictly convex* in p, that is no line segment with positive length passes through p and is contained in the boundary; the converse is not always true, as for instance the set

$$\{(x, y, z) \in \mathbb{R}^3 \mid x^4 + y^4 + z^4 \leqslant 1\}$$

is smooth and strictly convex but has vanishing Gaussian curvature in any extreme point along the x, y and z axis.

Now let ℓ_0 be an oriented line tangent to A. We equip $\mathbb{R}^3$ with a frame having ℓ_0 as z-axis, and parameterize $\mathcal{L}$, the space of oriented lines whose directions make a positive dot product with $(0, 0, 1)$ using the intersection points with the planes $z = 0$ and $z = 1$ as explained page 53.

Lemma 7.1. *If A is of class C^k and rotund in its contact point with ℓ_0 then there exists a neighborhood V of ℓ_0 in $\mathcal{L}$ such that $V \cap \partial\mathcal{T}(A)$ is C^{k-1}-diffeomorphic to $\mathbb{R}^3$.*

Proof. Let x_0 denote a point in the interior of A and S denote a sphere centered in x_0 of radius $r > 0$ sufficiently small to be also contained in the interior of A. Let U denote a neighborhood of p in ∂A such that A is of class C^k and rotund any point of U. Let π denote the central projection with center x_0 from U to S:

$$\pi(q) = x_0 + \frac{r}{\|x_0 q\|}\overrightarrow{x_0 q}.$$

Let ℓ be an oriented line tangent to A in $q \in U$. We let $\psi(\ell)$ denote the oriented line tangent to S in $\pi(q)$, contained in the plane spanned by x_0 and ℓ and whose direction vector makes an angle smaller than $\frac{\pi}{2}$ with that of ℓ (see Figure 39).

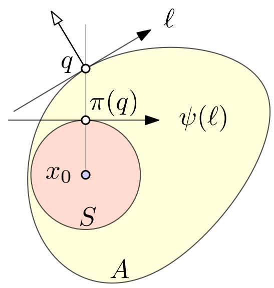

Figure 39: If A is of class C^k and rotund in its contact point with a tangent line ℓ_0 then $\partial\mathcal{T}(A)$ is a C^{k-1}-manifold near ℓ_0.

The map ψ is a C^{k-1}-diffeomorphism between the set of oriented lines tangent to A in a point of U and the set of oriented lines tangent to S in a point of $\pi(U)$. Indeed,

- the map π is of class C^k on U,

- the map associating to a point $q \in U$ the outward unit normal vector to A in q is of class C^{k-1} on U,
- the map that semds a line tangent to A in a point of U to that tangency point is of class C^{k-1} (this can be obtained,for instance, via properties of the gauge function associated to A and x_0).

The space of oriented lines tangent to S is diffeomorphic to the unit tangent bundle of the sphere $\mathbb{S}^2$, which is a smooth 3-dimensional manifold. □

Instead of working in the representation of $\mathcal{L}$ by $\mathbb{R}^4$ we could consider $\partial\mathcal{T}(A)$ a sub-manifold[3] of the Klein quadric G.a Lemma 7.1 remains valid as the mapping of $\mathcal{L}$, envisaged as a subset of G, to the representation by $\mathbb{R}^4$ is simply a central projection [PW01, Exemple 8.1.1] and induces a C^∞-diffeomorphism between (a subset of) G and $\mathbb{R}^4$. This change of point of view brings two differences. First, when A is of class C^k and rotund in all its boundary points the proof of Lemma 7.1 gives a *global* C^{k-1}-diffeomorphism between $\partial\mathcal{T}(A)$ and the unit tangent bundle of $\mathbb{S}^2$. Second, the tangent space in ℓ_0 to $\partial\mathcal{T}(A)$ no longer lives in the

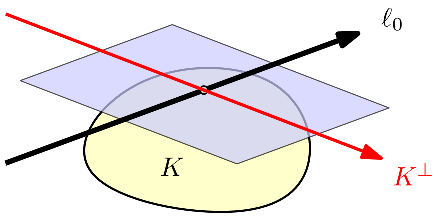

Figure 40: Definition of $A^{\perp\ell_0}$.

"space of lines": it is a $\mathbb{RP}^3$ that meets G in a two-dimensional section[4]. In contrast, working in the representation of $\mathcal{L}$ by $\mathbb{R}^4$ leads to a simple geometric interpretation of this tangent space. Let $A^{\perp\ell_0}$ denote the line tangent to A and perpendicular to ℓ_0 in its contact point with A; we orient $A^{\perp\ell_0}$ so that translating ℓ_0 in the direction of the outer normal to A in p makes it pass to the left of $A^{\perp\ell_0}$ (See Figure 40).

Lemma 7.2. *If A is rotund in its contact point with ℓ_0 then the tangent hyperplane to $\partial\mathcal{T}(A)$ in ℓ_0 is the subset of $\mathcal{L}$ of lines intersecting $A^{\perp\ell_0}$.*

Proof. Let p denote the contact point of A with ℓ_0 and r the inverse of the Gaussian curvature of A in p. Let B_1 and B_2 be

3 The differential structure on G can be obtained as follows: the pre-image of G through the quotient map $\mathbb{R}^6 \to \mathbb{RP}^5$ is a 5-dimensional linear cone $\tilde{G} \subseteq \mathbb{R}^6$, which is a smooth sub-manifold of $\mathbb{R}^6$; the quotient map then transports this differential structure from$\tilde{G}$ onto G.

4 Specifically, that section is the union of two $\mathbb{RP}^2$, which are the common line transversals to ℓ_0 and $A^\perp\ell_0$ (forming a line bundle and a line field).

two balls with respective radii $r_1 < r < r_2$ such that A, B_1 and B_2 are all internally tangent in p. There exists a neighborhood U of p such that:

$$B_1 \cap U \subseteq A \cap U \subseteq B_2 \cap U.$$

It follows that there exists a neighborhood V of ℓ_0 in $\mathcal{L}$ such that

$$\mathcal{T}(B_1) \cap V \subseteq \mathcal{T}(A) \cap U \subseteq \mathcal{T}(B_2) \cap U.$$

By Lemma 7.1, $\mathcal{T}(B_1)$, $\mathcal{T}(A)$ and $\mathcal{T}(B_2)$ are smooth in ℓ_0. These three sets must then be tangent in ℓ_0, and therefore have the same tangent hyperplane. Observe that $B_i^{\perp \ell_0}$ is the boundary of the screen of (B_i, ℓ_0), as defined page 31. Lemma 3.9 thus implies that the tangent hyperplane to $\mathcal{T}(B_i)$ in ℓ_0 is precisely the set of line transversals to $B_i^{\perp \ell_0}$. Since $B_1^{\perp \ell_0} = B_2^{\perp \ell_0} = A^{\perp \ell_0}$ the statement follows. $\square$

⋆ ⋆ ⋆

Now let $\mathcal{F} = \{A_1, \dots, A_n\}$ be a family of convex sets, each tangent to ℓ_0 in a rotund point. By Lemma 7.1 the sets $\partial\mathcal{T}(A_i)$ are, near ℓ_0, differentiable manifolds. With Lemma 7.2, the linear dependency of the normals to $\partial\mathcal{T}(A_i)$ in ℓ_0 follows from the classification of Chapter 6. Specifically, the normals to $\partial\mathcal{T}(A_1), \dots, \partial\mathcal{T}(A_4)$ are linearly dependent if and only if, up to relabelling the sets, one of the following configurations is realized (see Figure 41):

(i) two sets are tangent in a point of ℓ_0,

(ii) three sets meet in a point of ℓ_0,

(iii) three sets have a common support plane containing ℓ_0,

(iv) two sets meet in a point of ℓ_0 and two other sets have a common support plane containing ℓ_0,

(v) four of the sets are tangent to a hyperbolic paraboloid containing ℓ_0.

The linear independence of a quadruple of normals expresses the condition that the four corresponding manifolds intersect transversely in the origin. We therefore say that the family $\mathcal{F}$ is in *transverse position with respect to* ℓ_0 if it does not contain any sub-configuration of type (i)-(v).

Let H_i^+ denote the tangential cone to $\mathcal{T}(A_i)$ in the origin. We say that $\mathcal{F}$ *pins* ℓ_0 *at first order* if ℓ_0 is isolated in $\bigcap_i H_i^+$. Arguments

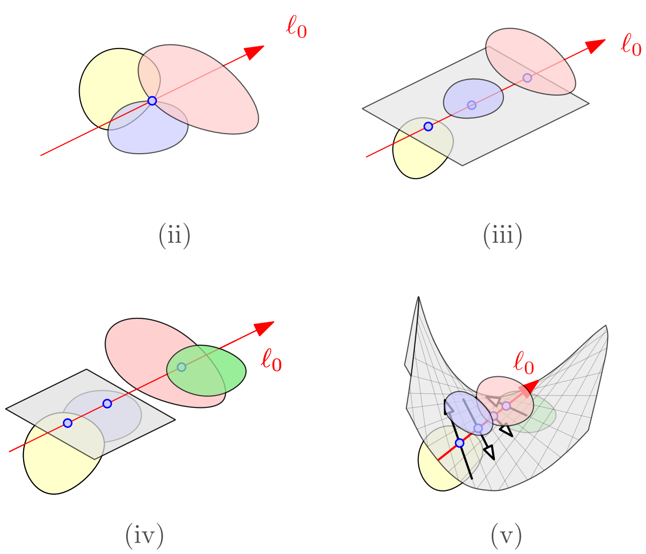

Figure 41: Configurations where the family is not in transverse position with respect to ℓ_0.

similar[5] to the proof of Lemma 6.1 show that if $\mathcal{F}$ pins ℓ_0 at first order then $\mathcal{F}$ pins ℓ_0 in the usual sense. The notion of first-order pinning is equivalent to the notion of stable pinning defined in the introduction of this chapter:

Lemma 7.3. *Let $\mathcal{F}$ be a family of convex sets, each tangent to one and the same line ℓ_0 in a rotund point. The family $\mathcal{F}$ is a stable pinning of ℓ_0 if and only if it pins ℓ_0 at first order.*

Proof. Let $\mathcal{F} = \{A_1, \ldots, A_n\}$ be a family of convex sets, each tangent to one and the same line ℓ_0 in a rotund point. Lemma 7.1 ensures that $\partial\mathcal{T}(A_i)$ is smooth in ℓ_0. Lemma 7.2 and Equation (5.2) imply that the outward normal $\eta_{A_i^{\perp\ell_0}}$ to $\mathcal{T}(A_i)$ in ℓ_0 writes $((1-\alpha)\vec{\eta}; \alpha\vec{\eta})$ where $\alpha \in \mathbb{R}$ describes the position of $A_i \cap \ell_0$ along ℓ_0, and $\vec{\eta}$ describes the normal to A_i in $A_i \cap \ell_0$ in the $\mathbb{S}^1$ of directions orthogonal to ℓ_0. It follows that applying a screw of axis ℓ_0 to A_i changes the normal $\eta_{A_i^{\perp\ell_0}}$ continuously in the parameters (angle of the rotation and amplitude of the translation) of the screw. As a consequence, if $\mathcal{F}$ pins ℓ_0 at first order, any sufficiently small perturbation of $\mathcal{F}$ pins ℓ_0 at first order, and any first-order pinning is stable.

To prove the converse we start by assuming that $\mathcal{F}$ does not pin at first order. All normals $\eta_{A_i^{\perp\ell_0}}$ are then contained in some closed halfspace H of $\mathbb{R}^4$ with the origin on its boundary. We argue that each $\eta_{A_i^{\perp\ell_0}}$ on the boundary of H can be perturbed

5 One need to be careful that in this setting, not being isolated does not guarantee the existence of a smooth path starting from ℓ_0 and moving inside $\bigcap_i \mathcal{T}(A_i)$; this is merely a technical gap, easily supplemented by standard compactness arguments.

so as to move in the interior of H. This step requires some care as the normals $\eta_{A_i^{\perp \ell_0}}$ live in a two-dimensional subspace Γ of $\mathbb{S}^3$; we show that not only does Γ not lie locally in any hyperplane (an argument already invoked in the proof of Theorem 3.10), but that it does not lie locally in any halfspace. $\square$

As a consequence, Steinitz's theorem (Theorem 5.3) ensures that any minimal stable pinning[6] of a line by convex sets tangent to the line in rotund points has size at most eight.

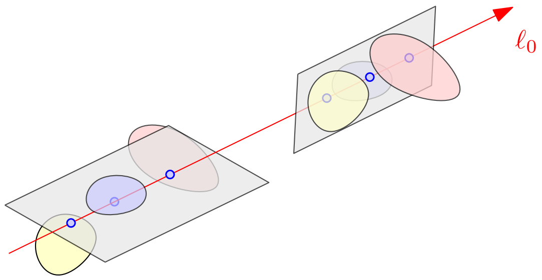

Figure 42: A stable pinning by convex sets that are not in transverse position.

Lemma 7.3 also implies that any pinning that is transverse is stable (and any minimal such pinning has size exactly five); the converse is not true, c.f. the example of Figure 42. In fact, Lemma 7.3 yields that the minimal stable pinnings of a line by convex sets tangent meeting it in rotund points are exactly the families $\{A_1, \ldots, A_n\}$ such that the family $\{A_1^{\perp \ell_0}, \ldots, A_n^{\perp \ell_0}\}$ is one of the configurations listed in Figure 38.

⋆ ⋆ ⋆

We finally restrict our attention to *semi-algebraic* ovaloids and obtain, under a mild non-degeneracy assumption, a pinning theorem. Let $\mathcal{F} = \{A_1, \ldots, A_n\}$ be a pinning of an oriented line ℓ_0 by a family of semi-algebraic convex sets, each of which is tangent to ℓ_0 in a rotund point.

As in Chapter 2, for $G \subseteq \mathcal{F}$ we let $\mathcal{K}(G) \subseteq \mathbb{S}^2$ denote the set of directions of line transversals to G. Recall that, as argued in the proof of Theorem 2.1, a family of convex sets in $\mathbb{R}^3$ has a line transversal with direction u if and only if any three members of the family has a line transversal with direction u; indeed, this is

6 Minimal stable pinning should be understood as "stable pinning of a line that does not properly contain another stable pinning of that line".

merely an application of Helly's theorem to the projection of the sets on some plane orthogonal to u. In other words,

$$\mathcal{K}(\mathcal{F}) = \bigcap_{T \in \binom{\mathcal{F}}{3}} \mathcal{K}(T) \tag{7.1}$$

where $\binom{\mathcal{F}}{3}$ denotes the set of triples of elements of $\mathcal{F}$. The map $G \to \mathbb{S}^2$ that associates to a line its direction induces a bijection between connected components of $\mathcal{T}(G)$ and connected components of $\mathcal{K}(G)$. It follows that Lemma 3.4 (iii) generalizes to convex sets [26, Lemma 3 (iv)], and a subfamily $G \subseteq \mathcal{F}$ pins ℓ_0 if and only if the direction of ℓ_0 is isolated in $\mathcal{K}(G)$. Combined with Identity (7.1), the assumption that $\mathcal{F}$ pins ℓ_0 implies that the direction $\overrightarrow{\ell_0}$ of ℓ_0 is an isolated point of the intersection of the $\mathcal{K}(T)$, where T ranges over the triples of elements in $\mathcal{F}$. We will argue that $\overrightarrow{\ell_0}$ is already isolated in the intersection of some at most four of those sets $\mathcal{K}(T)$.

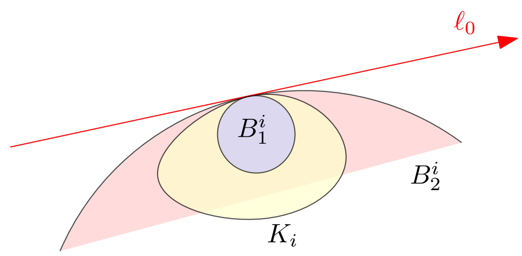

Figure 43: Sandwiching an ovaloid between two balls locally near its contact point with the line ℓ_0.

Let $T = \{A_a, A_b, A_c\}$ be some triple of elements in $\mathcal{F}$ and assume that $\overrightarrow{\ell_0} \in \partial\mathcal{K}(T)$. If T does not pin ℓ_0 and no two members of $\mathcal{T}$ are externally tangent in a point of ℓ_0 then $\partial\mathcal{K}(T)$ is smooth in $\overrightarrow{\ell_0}$. This statement can be proved using the explicit algebraic equations of the arcs of curves forming $\partial\mathcal{K}(T)$ when T consists of balls [26, Lemma 6]; the general case then follows by picking, for each ovaloid A_i, two balls B_1^i and B_2^i so that all three sets are internally tangent in $A_i \cap \ell_0$ and A_i is sandwiched, locally near $A_i \cap \ell_0$, between the two balls (see Figure 43).

We now argue that if no two members of $\mathcal{F}$ are externally tangent in a point of ℓ_0 then if $\mathcal{F}$ is a minimal pinning of ℓ_0 it must have cardinality at most twelve. Putting

$$C = \left\{ T \mid T \in \binom{\mathcal{F}}{3} \text{ and } \overrightarrow{\ell_0} \in \partial\mathcal{K}(T) \right\}$$

the condition that $\mathcal{F}$ pins ℓ_0 implies, together with Identity 7.1, that $\overrightarrow{\ell_0}$ is isolated in the intersection $\bigcap_{T \in C} \mathcal{K}(T)$. We can assume that no triple in $\mathcal{F}$ pins ℓ_0, as otherwise we are done. With the

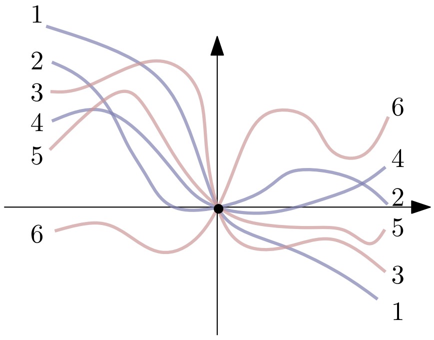

Figure 44: A "sandwich region" is locally defined by at most 4 curves in the plane: the fact that the origin is isolated among the points below the blue curves (1,2 and 4) and above the red curves (3, 5 and 6) can already be observed from the curves 1, 3 and 5.

condition that no two sets are externally tangent in a point of ℓ_0, we thus have that $\partial\mathcal{K}(T)$ is smooth in $\overrightarrow{\ell_0}$ for any $T \in C$. We can then recast $\mathcal{K}(\mathcal{F})$ near $\overrightarrow{u}$ as the region above the lower envelope and below the upper envelope of families of functions. These functions are semi-algebraic, so their upper and lower envelopes are defined, locally near $\overrightarrow{\ell_0}$, by two curves each (see Figure 44). We thus have that $\overrightarrow{\ell_0}$ is locally the single point above the lower envelope of two curves and below the upper envelope of two (possibly other) curves; the at most twelve objects in the union of the four triples of $\mathcal{F}$ defining these curves suffices to pin ℓ_0 [26, Theorem 12].

⋆ ⋆ ⋆

NOTES.

This chapter is based on a paper co-authored with Stefan König and Sylvain Petitjean [26] and an on-going work with Guillaume Batog.

The notion of "stable pinning" deviates slightly from the one introduced in Chapter 3 for families of balls: there, we allowed the balls to increase in radius. Increasing the radius of a ball B while keeping it tangent to ℓ_0 in the same point, with the same tangent plane at that point, leaves $B^{\perp \ell_0}$ unchanged and therefore does not affect the normal to $\mathcal{T}(B)$ in ℓ_0. These notions of perturbations are thus equivalent for our purpose (and even more general perturbation schemes also lead to a notion of stability still equivalent to first-order pinning).

The condition of Lemma 7.1, that the convex A touches ℓ_0 in a rotund point, can be weakened: indeed, it suffices that A has positive curvature in the plane defined by ℓ and the normal to A in the contact point with ℓ_0. This implies, for instance, that if A is a cylinder tangent to but not containing ℓ_0 then $\partial\mathcal{T}(A)$ is a smooth manifold near ℓ_0. If, however, A is not C^2 in its contact point with ℓ_0 then $\partial\mathcal{T}(A)$ may not be a differentiable manifold; for instance, if A is a polytope then $\partial\mathcal{T}(A)$ is smooth only in the neighborhood of a line ℓ_0 tangent to A in a single point, interior to an edge. For arbitrary convex sets A, we proved that $\partial\mathcal{T}(A)$ is always a *topological* manifold.

The condition that four sets are tangent to a parabolic hyperboloid along one of its rulings amounts to a projective relation between the tangency points of the sets with that rulings and their tangent planes at those points. Indeed, if $\mathcal{P}$ is a parabolic hyperboloid and Δ a line contained in $\mathcal{P}$ then the map sending a point $x \in \Delta$ to the plane tangent to $\mathcal{P}$ in x is a projective transformation [PW01, Theorem 3.2.9].

Other questions on sets of lines tangent to convex sets may benefit from the differential geometry perspective unfolded here. For instance, the question of characterizing quadruples of objects in $\mathbb{R}^3$ with infinitely many common tangents received some attention lately as such families induce degeneracies in global visibility structure. When the objects are balls or lines, this question was answered by closely inspecting algebraic systems describing these common tangents [6, MPT01, Meg01, The02, MST03, MS05, ST08]; this approach, however, does not extend easily even to the case of four ellipsoids in $\mathbb{R}^3$. A necessary condition for a quadruple of smooth rotund convex sets to have infinitely many common tangents is that they are in non-transverse position with respect to all of their common tangents, considerably narrowing down the range of configurations to analyze.

Part III

HELLY NUMBERS AND SIMPLICIAL COMPLEXES

8 AROUND THE GRÜNBAUM-MOTZKIN CONJECTURE

Given two families $\mathcal{F}$ and G of sets and an integer $r \geqslant 1$, we say that $\mathcal{F}$ is a (G, r)*-family* if the intersection of any subfamily of $\mathcal{F}$ is a disjoint union of at most r members of G. Call a family of sets *non-additive* if any union of disjoint members of the family is *not* in the family. In 1961, Grünbaum and Motzkin [GM61] formulated the following conjecture:

Conjecture 8.1 (Grünbaum-Motzkin)**.** *If $\mathcal{F}$ is a (G, r)-family where G is non-additive and closed under intersection then the Helly number of $\mathcal{F}$ is at most r times the Helly number of G.*

A proof of this conjecture was published by Eckhoff and Nischke [EN09], building on ideas of Morris [Mor73]. The conjecture was previously settled in two special cases: when G is a family of compact convex subsets in $\mathbb{R}^d$, by Amenta [Ame96], and when G is a good cover[1] in $\mathbb{R}^d$, by Kalai and Meshulam [KM08]. These three proofs seem, at first glance, rather different: Eckhoff and Nischke use a generalized pigeonhole principle, Amenta uses ideas from combinatorial optimization and Kalai and Meshulam use techniques from homology theory. In this chapter, we show that these proofs share a similar core, a theorem on projections of a simplicial complex.

NOTATIONS. For $n \in \mathbb{N}$ we let $[n]$ denote the set $\{0, \ldots, n\}$. We shorten $\bigcup_{x \in F} x$ and $\bigcap_{x \in F} x$ into, respectively, $\bigcup_F$ and $\bigcap_F$ and denote by 2^F the set $\{Z \mid Z \subseteq F\}$ of subsets of F (including the empty set).

⋆ ⋆ ⋆

At the combinatorial level, a *simplicial complex* X over a set of *vertices* V is a non-empty family of subsets of V closed under taking subsets; in particular, $\emptyset$ belongs to every simplicial complex. An element σ of X is a *simplex*; its *dimension* is the cardinality of σ minus one. A simplex τ contained in a simplex σ is a *face* of σ; if $\dim \tau = \dim \sigma - 1$ we say that τ is a *facet* of σ. The k-dimensional *skeleton* of a simplicial complex X is the set of simplices of X of dimension at most k (and is also a simplicial complex). Given a

1 Recall that a good cover is a family of open sets such that the intersection of any subfamily is empty or contractible.

subset $S \subseteq V$, the sub-complex of X *induced* by S, denoted $X[S]$, is the set of simplices of X contained in S; it is easily seen to be a simplicial complex.

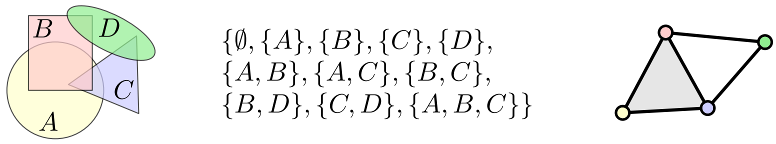

Figure 45: A family of subsets of $\mathbb{R}^2$ (left), its nerve (middle) and a geometric realization of that nerve (right).

To any finite family $\mathcal{F}$ of sets is associated the simplicial complex

$$N(\mathcal{F}) = \{H \subseteq \mathcal{F} \mid \cap_H \neq \emptyset\},$$

called the *nerve* of $\mathcal{F}$ (see Figure 45). We define a *simplicial hole* of a simplicial complex X with vertex set V to be a subset $S \subseteq V$ such that $X[S] = 2^S \setminus \{S\}$ (see Figure 46); if S is a simplicial hole of X then $X[S]$ is (the face lattice of) the boundary of a simplex, hence the name. With these definitions, the Helly number of a family reformulates as:

Claim 8.2. *The Helly number of a family $\mathcal{F}$ is the maximum cardinality of a simplicial hole of its nerve $N(\mathcal{F})$.*

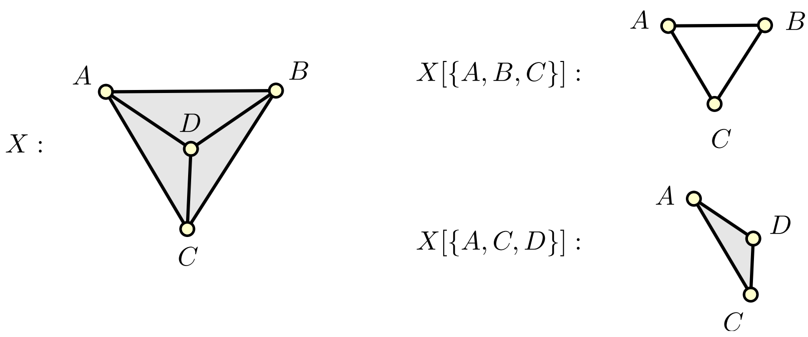

Figure 46: A (geometric representation of a) nerve (left) and two induced subcomplex, one that is a simplicial hole (top-right) and one that is not (bottom-right).

Now, back to the conjecture of Grünbaum and Motzkin. The motivation for requiring that the "ground" family G be non-additive and closed under intersection is the following property [GM61, Theorem 1]:

Lemma 8.3. *Let D_1 and D_2 be two subfamilies of a family G that is non-additive and closed under intersection. If $\bigcup_{D_1} = \bigcup_{D_2}$ and each D_i is a disjoint family then $D_1 = D_2$.*

Proof. Let $A \in D_1$ and write $A = \bigcup_{Z \in D_2} A \cap Z$. Since G is closed under intersection and non-additive, there can be only one non-empty term in the right-hand side: there exists $A' \in D_2$ such that $A \subseteq A'$. Writing $A' = \bigcup_{Z \in D_1} A' \cap Z$, the same argument yields that at most one term in the union is non-empty, and that must be $A' \cap A$; it follows that $A = A'$. Applying this argument to every A in D_1 we get that $D_1 \subseteq D_2$, and exchanging the roles of D_1 and D_2 concludes the proof. □

If G is non-additive and closed under intersection then any element A in a (G, r)-family decomposes *uniquely* into a disjoint union $A = Z^1 \cup \ldots \cup Z^t$ of at most r members of G; we call the Z^i the *components* of A (over G).

For the sake of the exposition, assume that no two elements of $\mathcal{F}$ have a component in common. Let $\mathcal{F}$ be a (G, r)-family and let C denote the union, for all elements $A \in \mathcal{F}$, of the set of components of A. Let $\pi : C \to \mathcal{F}$ be the function that maps each element of C to the element of $\mathcal{F}$ it is a component of. We extend π to a map from 2^C to $2^{\mathcal{F}}$ by putting

$$\pi(\{Z_1, \ldots, Z_k\}) = \{\pi(Z_1), \ldots, \pi(Z_k)\},$$

making π a *simplicial map*. We call a simplicial map *dimension-preserving* if any simplex is mapped to a simplex of the same dimension. We define the *multiplicity* of a simplicial map as the maximum number of simplices that are mapped to the same simplex; in other words, the multiplicity is the maximum cardinality of a fiber. We also call a map *at most* r*-to-one* if it has multiplicity at most r.

Lemma 8.4. *If $\mathcal{F}$ is a (G, r)-family and G is non-additive and closed under intersection then π induces a surjective, dimension-preserving simplicial map at most r-to-one from the nerve of C to the nerve of $\mathcal{F}$.*

Proof. Two components of the same set are disjoint so if $\sigma = \{Z_1, \ldots, Z_t\}$ is a simplex of $N(C)$ then the $\pi(Z_i)$ are pairwise distinct, and $\pi(\sigma)$ has the same cardinality as σ. The restriction of π to $N(C)$ is thus dimension-preserving. For any simplex τ in $N(\mathcal{F})$ we have that $\bigcap_\tau$ is a disjoint union of the $\bigcap_\sigma$ where σ ranges over $\pi^{-1}(\tau)$; thus, $\pi(N(C)) = N(F)$ and the fact that $\mathcal{F}$ is a (G, r)-family implies that $\pi^{-1}(\tau)$ has cardinality at most r. □

We can now outline the common reformulation of the proofs of Kalai-Meshulam, Amenta and Eckhoff-Nischke in a two-step argument most apparent in the presentation of Kalai-Meshulam. Let C denote the set of components of members of $\mathcal{F}$, where "component" is understood either relatively to a family G (as above) or in the topological sense of connected component. The first step is to define some index $\rho(X)$ of a simplicial complex X such that $\rho(N(C))$ can be bounded and the Helly number of $\mathcal{F}$ can be controlled in terms of $\rho(N(\mathcal{F}))$. The second step is to study how this index ρ behaves under a dimension-preserving simplicial map $\phi : X \to Y$ between two simplicial complexes.

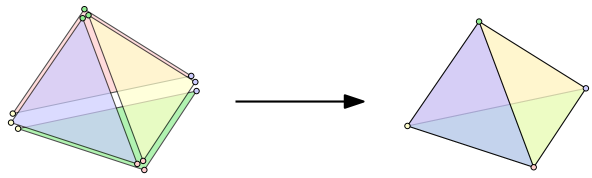

Figure 47: A map of multiplicity larger than 1 can create simplicial holes.

Just like ϕ may glue independent simplices of X into a simplicial hole of Y (see Figure 47), the value of $\rho(Y)$ cannot be bounded solely as a function of $\rho(X)$. Remarkably, in all three cases this can be taken care of by taking into account the multiplicity of ϕ.

⋆ ⋆ ⋆

Kalai and Meshulam [KM08] proved a bound on the Helly number of a (G, r)-family where G is a good cover in $\mathbb{R}^d$. They use techniques from homology theory[2], a standard approach to formalizing the notion of "hole" in a topological space or a simplicial complex. The *Leray number* $L(X)$ of a simplicial complex X, is defined as

$$L(X) = \min\{i \in \mathbb{N} \mid \forall S \subseteq V, \forall j \geqslant i, \quad \tilde{H}_j(X[S]) = 0\},$$

where V is the set of vertices of X and $\tilde{H}_j(Y)$ denotes the j-dimensional reduced $\mathbb{Q}$-homology of Y. If Y a simplicial hole with k vertices then $\tilde{H}_{k-2}(Y) \neq 0$ and $L(Y)$ is therefore at least $k-1$. In other words, the maximum cardinality of a simplicial hole exceeds the Leray number by at most one; the Helly number of $\mathcal{F}$ is thus at most $L(N(\mathcal{F})) + 1$.

The first step in the proof of Kalai and Meshulam is a projection theorem [KM08, Theorem 1.3] that reformulates as follows:

Theorem 8.5. *Let* $\pi : X \to Y$ *be a simplicial map of multiplicity* r *between two simplicial complexes. If* π *is surjective and dimension-preserving then* $L(Y) + 1 \leqslant r(L(X) + 1)$.

2 We refer to Appendix B for a brief overview of the notions of homology and homotopy we use.

Theorem 8.5 is proven via a rather sophisticated machinery: a homology spectral sequence, constructed by Goryunov and Mond [GM93], that computes the homology of the projection of a space.

Now, let G be a good cover in $\mathbb{R}^d$ and let $\mathcal{F}$ be a (G, r)-family. Let C be the set of components of elements of $\mathcal{F}$ over G and let $\pi : N(C) \to N(\mathcal{F})$ denote the simplicial map induced by the relation "being a component of". By Lemma 8.4, π is dimension-preserving and has multiplicity at most r. Thus, the Helly number of $\mathcal{F}$ is at most $L(N(\mathcal{F})) + 1$ which is at most $rL(N(C)) + r$ by Theorem 8.5. It remains to bound the Leray number of $N(C)$; this is done using the classical *Nerve Theorem* of Borsuk:

Lemma 8.6. *If C is a good cover in $\mathbb{R}^d$ then $L(N(C))$ is at most d.*

Proof. The nerve theorem asserts that $N(C)$ is homotopy-equivalent to $\bigcup_C$. Now, an open subset of $\mathbb{R}^d$ has trivial homology in dimension d or larger [Gre67, p. 121], and homology is preserved under homotopy. It follows that $N(C)$ has trivial homology in any dimension $j \geqslant d$. The same argument holds for $N(S)$ where S is any subset of C; since $N(S) = N(C)[S]$ it follows that the Leray number of $N(C)$ is at most d. □

With the previous argument, this yields a bound of $r(d+1)$ on the Helly number of $\mathcal{F}$. Note that this does not prove that the Helly number of $\mathcal{F}$ is at most r times that of G. In general there can be an arbitrarily large gap between the Leray number of a simplicial complex and the maximum dimension of a simplicial hole (consider, for instance, the barycentric subdivision of a simplicial hole with k vertices).

⋆ ⋆ ⋆

Amenta [Ame96] proved a bound on the Helly number of (G, r)-families, where G is a family of convex compact sets in $\mathbb{R}^d$. While this is essentially[3] a special case of the situation where G is a good cover, her proof uses a different projection theorem that avoids the use of the spectral sequence artillery. We reformulate this proof in the language of simplicial complexes.

We call a map $\alpha : 2^{[n]} \to \overline{\mathbb{N}}$ a *good graduation* if for any subsets $\sigma, \tau \subseteq [n]$ we have:

- $\sigma \subset \tau \Rightarrow \alpha(\sigma) \leqslant \alpha(\tau)$, and
- $\alpha(\sigma) = \alpha(\tau) \Rightarrow \alpha(\sigma \cup \tau) = \alpha(\sigma) = \alpha(\tau)$.

3 Up to the distinction open/compact, c.f. the *Notes* section.

These conditions ensure that for any $i \in \mathbb{N}$, the set $\alpha^{-1}([i])$ is a simplicial complex and $\alpha^{-1}([i+1]) \setminus \alpha^{-1}([i])$ is empty or contains a unique inclusion-wise maximal simplex. If X is a simplicial complex with vertex set $[n]$ and α is a good graduation such that $X = \alpha^{-1}(\mathbb{N})$ then we call α a *good graduation of* X; in that case, α induces a filtration

$$\emptyset \subseteq \alpha^{-1}([0]) \subseteq \alpha^{-1}([1]) \subseteq \ldots \subseteq \alpha^{-1}(\mathbb{N}) = X$$

of X by a sequence of simplicial complexes where two distinct consecutive complexes differ only in the addition of a simplex and all its missing faces. A *simplicial hole of* α is a simplicial hole of some $\alpha^{-1}([i])$.

Lemma 8.7. *Let α be a good graduation. If H is a simplicial hole of α of cardinality at least 2 then a facet of H is also a simplicial hole of α.*

Proof. Since H is a simplicial hole of α, no two of its facets have the same image under α. Let H' be the facet of H whose image under α is largest. Any facet σ of H' writes $\sigma = H \cap \tau$ where τ is a facet of H, and therefore $\alpha(\sigma) \leqslant \alpha(\tau) < \alpha(H')$. It follows that H' is a simplicial hole of α. □

Call a simplicial hole H of α *temporary* if $\alpha(H) \in \mathbb{N}$, i.e. if H is not a simplicial hole of $\alpha^{-1}(\mathbb{N})$, and let $\Delta(\alpha)$ denote the maximum cardinality of a temporary simplicial hole of α. Lemma 8.7 implies that the Helly number of a family $\mathcal{F}$ is at most $\Delta(\alpha)+1$, where α is any good graduation of the nerve of $\mathcal{F}$.

Let $\pi : X \to Y$ be a simplicial map of multiplicity r between two simplicial complexes with respective vertex sets $[x]$ and $[y]$. To any map $\alpha : 2^{[x]} \to \overline{\mathbb{N}}$ we associate

$$\beta : \begin{cases} 2^{[y]} & \to \quad \overline{\mathbb{N}} \\ \sigma & \mapsto \quad \begin{cases} \min_{\pi^{-1}(\sigma)} \alpha \text{ if } \sigma \in \pi(X) \\ \infty \text{ if } \sigma \in 2^{[y]} \setminus \pi(X). \end{cases} \end{cases} , \tag{8.1}$$

and we let $\rho : Y \to X$ denote any section of π satisfying $\beta = \alpha \circ \rho$.

Lemma 8.8. *If α is a good graduation of X then β is a good graduation of Y.*

Proof. It is straightforward to verify that if $\sigma \subset \tau \in Y$ then $\beta(\sigma) \leqslant \beta(\tau)$. If $\beta(\sigma) = \beta(\tau) = i$ then $\alpha(\rho(\sigma)) = \alpha(\rho(\tau)) = i$. Since α is a good graduation this implies that $\alpha(\rho(\sigma) \cup \rho(\tau)) = i$. Now, π is a simplicial map, so $\pi(\rho(\sigma) \cup \rho(\tau)) = \pi(\rho(\sigma)) \cup \pi(\rho(\tau)) = \sigma \cup \tau$. It follows that $\beta(\sigma \cup \tau) = i$, which completes the proof. □

The key argument in the proof of Amenta is the following projection theorem:

Lemma 8.9. *If the map* π *is surjective and dimension-preserving then* $\Delta(\beta)+1 \leqslant r(\Delta(\alpha)+1)$.

Proof. Let H be a temporary simplicial hole of β. We bound the dimension of H by constructing two sequences H_i and S_i with the following properties:

(i) $\forall v \in H_i \setminus S_i$, $\beta(H_i \setminus \{v\}) < \beta(H_i)$.

(ii) $\forall j \leqslant i$ there exists $\tau \in \pi^{-1}(H_i)$ such that $\alpha(\tau) = \beta(H_j)$.

For any $\sigma \in X$ we let $B(\sigma)$ denote an inclusionwise-minimal subset of σ such that $\alpha(B(\sigma)) = \alpha(\sigma)$. The minimality of $B(\sigma)$ ensures that it is a simplicial hole of α and has therefore cardinality at most $\Delta(\alpha)$. We construct (H_i) and (S_i) by induction, starting with $H_1 = H$ and $S_1 = \pi(B(H_1))$; Property (i) holds because H is a simplicial hole and property (ii) can be seen to hold by taking $\tau = \rho(H_1)$.

Assume that the sequences $H_1, \ldots, H_i$ and $S_1, \ldots, S_i$ exist and satisfy (i) and (ii). Since β is a good graduation, property (i) implies that for any two distinct $v, v' \in H_i \setminus S_i$ we have $\beta(H_i \setminus \{v\}) \neq \beta(H_i \setminus \{v'\})$. Let v_i denote the element in $H_i \setminus S_i$ such that $\beta(H_i \setminus \{v_i\})$ is maximal. We define $H_{i+1} = H_i \setminus \{v_i\}$ and $S_{i+1} = S_i \cup \pi(B(\rho(H_{i+1})))$. For any $v \in H_{i+1} \setminus S_i$ we have:

$$\beta(H_{i+1} \setminus \{v\}) \leqslant \beta(H_i \setminus \{v\}) < \beta(H_i \setminus \{v_i\}) = \beta(H_{i+1}).$$

Since $H_{i+1} \setminus S_{i+1} \subseteq H_{i+1} \setminus S_i$, the pair (H_{i+1}, S_{i+1}) satisfies property (i). It remains to check that (H_{i+1}, S_{i+1}) satisfies property (ii). For $j = i+1$ we can simply take $\tau = \rho(H_{i+1})$ to have $\alpha(\tau) = \beta(H_{i+1})$. Let $j \leqslant i$. From

$$\pi(B(\rho(H_j))) \subseteq H_{i+1} \subseteq H_j = \pi(\rho(H_j)),$$

we get that there exists τ such that $B(\rho(H_j)) \subseteq \tau \subseteq \rho(H_j)$ and $\pi(\tau) = H_{i+1}$. Then, since

$$\alpha(B(\rho(H_j))) \leqslant \alpha(\tau) \leqslant \alpha(\rho(H_j)),$$

we have that $\alpha(\tau) = \alpha(\rho(H_j)) = \beta(H_j)$ and property (ii) also holds for (H_{i+1}, S_{i+1}) and any $j \leqslant i$.

We iterate this construction until $H_i = S_i$. A simple induction shows that $|H_i| = |H| - i + 1$ and $|S_i| \leqslant i\Delta(\alpha)$, so if the construction stops after m steps we have:

$$|H| \leqslant i\Delta(\alpha) + m - 1.$$

Property (i) ensures that for any $i < j$ we have $\beta(H_i) > \beta(H_j)$ and property (ii) forces $m \leqslant \#\pi^{-1}(H_m) \leqslant r$. Altogether, we obtain that H has cardinality at most $r\Delta(\alpha) + r - 1$. Since this holds for any temporary simplicial hole H of β, the statement follows. $\square$

We can now conclude the proof of Amenta's theorem. Let G be a family of compact convex sets in $\mathbb{R}^d$, let $\mathcal{F}$ be a (G, r)-family and let C denote the set of connected components of members of $\mathcal{F}$ (equivalently, C is the set of components of $\mathcal{F}$ relative to G). We equip $\mathbb{R}^d$ with the lexicographic ordering $\prec$ and for every simplex $\sigma \in N(C)$ let p_σ denote the point of $\bigcap_\sigma$ smallest with respect to $\prec$. We let $p_{-\infty}$ denote a point smaller than all the p_σ's, collect all these points in a set $P = \{p_{-\infty}\} \cup \{p_\sigma \mid \sigma \in N(C)\}$ and number them in increasing order:

$$P = \{p_1, \dots p_m\} \quad \text{with} \quad p_i \prec p_{i+1} \quad \text{for} \quad 1 \leqslant i \leqslant n-1.$$

Let $\alpha : 2^{[n]} \to \overline{\mathbb{N}}$ map each simplex $\sigma \in N(C)$ to the integer i such that $p_\sigma = p_i$ and each simplex $\sigma \notin N(C)$ to $+\infty$. We define β as in Equation 8.1.

Lemma 8.10. *The maps α and β are good graduations of, respectively, $N(C)$ and $N(\mathcal{F})$ and $\Delta(\alpha) \leqslant d$.*

Proof. It is clear that α is monotone and $\alpha^{-1}(\{i\})$ has a unique maximal face, namely the set of all elements in C containing p_i. Moreover, $\alpha^{-1}(\mathbb{N}) = N(C)$ and $\beta^{-1}(\mathbb{N}) = N(F)$. Thus, α is a good graduation of $N(C)$, and, by Lemma 8.9, β is a good graduation of $N(\mathcal{F})$. Let H be a temporary simplicial hole of α and $i = \alpha(H)$ and let $A = \{p \mid p \in \mathbb{R}^d \text{ and } p \preceq p_{i-1}\}$. Since $(\bigcap_H) \cap A$ is empty and A is convex, Helly's theorem asserts that some subfamily $S \subseteq H \cup \{A\}$ of cardinality at most $d+1$ has empty intersection. Since H is a temporary simplicial hole, for any facet σ of H we have $\alpha(\sigma) \leqslant i-1$ and thus $(\bigcap_\sigma) \cap A$ is nonempty. It follows that S must contain H. Since $\bigcap_H$ is nonempty, S must also contain A, and H has cardinality at most d. □

We deduce from Lemma 8.7 that the Helly number of $\mathcal{F}$ is at most $\Delta(\beta)+1$. By Theorem 8.9, $\Delta(\beta)$ is at most $r\Delta(\alpha)+r-1$ and by Lemma 8.10, $\Delta(\alpha)$ is at most d. Putting everything together we obtain that the Helly number of $\mathcal{F}$ is at most $r(d+1)$.

⋆ ⋆ ⋆

We finally turn our attention to the proof of the Grünbaum-Motzkin conjecture by Eckhoff and Nischke [EN09]. The reformulation of that proof in the language of simplicial complexes reveals a statement that is surprisingly general.

The proof by Eckhoff and Nischke is formulated in the combinatorial language of independent families. An *independent family* of a set X is a family of pairwise disjoint subsets of X. The *meet* of two independent families P and Q is defined as:

$$P \wedge Q = \{A \cap B \mid A \in P, B \in Q \text{ and } A \cap B \neq \emptyset\}.$$

The meet of two independent families is again an independent family. From

$$P_1 \wedge P_2 = P_2 \wedge P_1 \quad \text{and} \quad (P_1 \wedge P_2) \wedge P_3 = P_1 \wedge (P_2 \wedge P_3)$$

it follows that we can define the meet of three or more independent families inductively without ambiguity. Given a family $\mathcal{P} = \{P_1, \dots, P_n\}$ of independent families, we define

$$\gamma(\mathcal{P}) = \max_{\mathcal{G} \subset \mathcal{P}} |\wedge_{P_i \in \mathcal{G}} P_i|,$$

that is, the maximum number of non-empty sets in the meet of a subfamily of $\mathcal{P}$. The key ingredient in the proof of Eckhoff and Nischke is the following pigeonhole principle (proven by direct combinatorial arguments).

Lemma 8.11. *Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a family of independent families of a finite set L. If $\left|L \setminus \bigcup_{P_i}\right| \leqslant 1$ for every i then there exists a subset $M \subseteq L$ with $|M| \geqslant \frac{|L|-1}{\gamma(\mathcal{P})} + 1$ such that every P_i contains an element C_i with $|M \setminus C_i| \leqslant 1$.*

The arguments used by Eckhoff and Nischke to deduce the Grünbaum-Motzkin conjecture from Lemma 8.11 reformulate into the following lifting theorem:

Theorem 8.12. *Let $\pi : X \to Y$ be a dimension-preserving, surjective simplicial map of multiplicity r between two simplicial complexes. If Y contains the $(k-1)$-skeleton of the k-simplex then X contains the $(\delta - 1)$-skeleton of the k-simplex, where $\delta = \lfloor \frac{k}{r} \rfloor$.*

Proof. Assume that $B \subseteq Y$ is the $(k-1)$-skeleton of a k-simplex and let $A = \pi^{-1}(B)$. We denote by $v_1, \dots, v_{k+1}$ the vertices of B. For every $(k-1)$-simplex σ of B we choose a simplex $\rho(\sigma) \in \pi^{-1}(\sigma)$ and collect all these pre-images in a set $L = \{\rho(\sigma) \mid \sigma \in B \text{ and } \dim \sigma = k-1\}$.

We first define a family of independent families of L. To every vertex $w \in A$ we associate the (possibly empty) family $D_w = \{\tau \mid w \in \tau \in L\}$. Then for v_i of B we put the independent family

$$P_i = \{D_w \mid D_w \neq \emptyset \text{ and } w \in \pi^{-1}(v_i)\}.$$

For every vertex v_i of B there is a unique $(d-1)$-simplex of B that does not contain v_i, and thus a unique element of L that does not contain a vertex of $\pi^{-1}(v_i)$; it follows that $\left|L \setminus \bigcup_{P_i}\right| \leqslant 1$ for every $1 \leqslant i \leqslant k+1$. Let $\mathcal{P} = \{P_1, \dots, P_{k+1}\}$.

We next argue that $\gamma(\mathcal{P})$ is at most r. Let $\mathcal{G} = \{P_{i_1}, \dots, P_{i_t}\} \subseteq \mathcal{P}$ and put $V = \{v_{i_1}, \dots, v_{i_t}\}$. We let $U = \{\{v_{i_1}, \dots, v_{i_t}\} \mid \forall 1 \leqslant j \leqslant$

$t, w_{i_j} \in \pi^{-1}(v_{i_j})\}$ denote the set of all possible lifts of V through π^{-1}. Since the sets in each D_w are pairwise disjoint, we have:

$$|V_{\mathcal{G}}| = \left|\left\{ \bigcap_{w \in W} D_w \mid W \in U \text{ and } \bigcap_{w \in W} D_w \neq \emptyset \right\}\right|.$$

Given $W \in U$, the intersection $\bigcap_{w \in W} D_w$ is nonempty if and only if there exists a simplex $\sigma \in X$ that contains W. This can only happen if W is a simplex of A. Since π maps any element of U that is a simplex of A to the same simplex of B, namely V, it follows that at most r elements of U can be simplices of A. We thus have $|V_{\mathcal{G}}| \leqslant r$ and the claim follows.

Now, let $\delta = \lfloor \frac{|L|-1}{\gamma(\mathcal{P})} \rfloor = \lfloor \frac{k}{r} \rfloor$. By Lemma 8.11 there exists $M \subseteq L$ with $|M| \geqslant \delta + 1$ such that every P_i contains an element D_i with $|M \setminus D_i| \leqslant 1$. We let w_i be the vertex in $\pi^{-1}(v_i)$ such that $D_{w_i} = D_i$ and let $S = \{w_1, \ldots, w_{k+1}\}$. Each D_{w_i} contains M except for at most one element. It follows that the intersection of λ of the D_{w_i} has size at least $|M| - \lambda \geqslant \delta + 1 - \lambda$. In particular any intersection of δ of the D_{w_i} is nonempty. It means that for any subset $\sigma \subset S$ of cardinality δ there exists some $\tau \in L$ such that $\sigma \subset \tau$; in particular, since τ is in A, so is σ. This proves not only that the $(\delta - 1)$-skeleton of the d-simplex S is contained in A, but that the union of *any* choice of lifts of the facets of B contains the announced skeleton. □

Now, assume that $\mathcal{F}$ is a (G, r)-family where G is non-additive and closed under intersection. Let C denote the family of components of members of $\mathcal{F}$ relatively to G and let $\pi : N(C) \to N(\mathcal{F})$ be the simplicial map induced by the relation "being a component of". By Lemma 8.4, π is dimension-preserving and has multiplicity at most r. Let h denote the Helly number of C and assume, for the contradiction, that the nerve of $\mathcal{F}$ has a simplicial hole S of cardinality $rh + 1$. This means that $N(\mathcal{F})$ contains the $(rh - 1)$-skeleton of the (rh)-simplex S. By Theorem 8.12, $N(C)$ must contain the h-skeleton of a (rh)-simplex T and $N(C)$ therefore contains the simplex T. The projection of the vertices of T are the vertices of S, so $\pi(T) = S$. It follows that $N(\mathcal{F})$ contains the whole simplex S, contradicting the initial assumption that S was a simplicial hole of $N(\mathcal{F})$.

⋆ ⋆ ⋆

NOTES.

This translation in simplicial terms of the Grünbaum-Motzkin conjecture is part of a joint work with Éric Colin de Verdière and

Grégory Ginot [19] and was inspired by the proof of Kalai and Meshulam [KM08].

The statements of Eckhoff-Nischke, Kalai-Meshulam, and Amenta on Helly numbers essentially generalize one another. It should be noted that the same does not seem to hold for the projection theorems (Theorems 8.12, 8.5 and 8.9).

Kalai and Meshulam state Lemma 8.6 in the case of *compact* good cover: families of compact sets where the intersection of every subfamily is empty or contractible. In their proof, they use the property that the union of the good cover in $\mathbb{R}^d$ has trivial homology in dimension $j \geqslant d$. As we mentioned, this is true if the sets are open. If they are compact, however, one should be cautious as already in $\mathbb{R}^2$ compact sets may have nontrivial homology in arbitrary high dimension (c.f. the so-called *Hawaiian earrings* [MB62]). This step is not detailed by Kalai and Meshulam; we note that assuming that the union of the sets admits a finite triangulation takes care of the matter [19, Lemma 21].

Amenta's proof uses the assumption that the sets in G are compact in two places. First, the compactness ensures that the family G is non-additive; the same holds if the sets are open (but *some* condition is needed as general convex sets do not make a non-additive family). Second, the compactness is used when p_σ is defined as the minimum point of $\bigcap_\sigma$. This is merely a convenient way to define the filtration, and compactness is not necessary.

Let G be a family of sets that is non-additive and closed under intersection. Grünbaum and Motzkin observed [GM61, Theorem 2] that if the intersection of any at most r members of a family $\mathcal{F}$ is a disjoint union of at most r members of G then $\mathcal{F}$ is a (G, r)-family. The number of components of intersections needs therefore only be checked for families of small cardinality.

Amenta's sweep argument is reminiscent of a result of Wegner [Weg75] from 1975. Wegner showed that the nerve of any finite collection of compact convex sets is d-*collapsible*, i.e. can be transformed into the empty set by a sequence of *collapses* and deletions of maximal simplices of dimension at most $d-1$. Considered backward a d-collapse is a filtration, where every intermediate structure is a simplicial complex. Wegner's filtration is obtained by sweeping the space with a hyperplane, and his bound on the dimension of the deleted simplices is similar to the proof of Lemma 8.10. Wegner conjectured that his result (the d-collapsibility of nerves) holds not only for convex sets but also for good covers; this was recently disproved by Tancer [Tan10].

A *LP-type* problem is a pair (Γ, w) where Γ is a set (the constraints) and w a function from 2^Γ to some totally ordered set, say $\overline{\mathbb{N}}$, satisfying two conditions:

- $\sigma \subset \tau \Rightarrow w(\sigma) \leqslant w(\tau)$, and
- if $\sigma \subset \tau$ and $w(\sigma) = w(\tau)$ then for any $\gamma \in \Gamma$ we have: $w(\sigma) < w(\sigma \cup \{\gamma\}) \Leftrightarrow w(\tau) < w(\tau \cup \{\gamma\})$.

A LP-type problem defines a filtration of 2^Γ by the sequence $(w([i])_{i \in \overline{\mathbb{N}}}$, each of which is a simplicial complex by the first condition. One can define a (temporary) simplicial hole of w as a set H such that $w(H) > w(\sigma)$ for any face σ of H (and $w(H) \in \mathbb{N}$). In the language of LP-type problems, a simplicial hole is called a *basis*, a temporary simplicial hole is called a *feasible basis*, and the index $\Delta(w)$ is the *combinatorial dimension* of (γ, w).

LP-type problems were introduced as combinatorial abstractions of the *linear programming* optimization problem by Sharir and Welzl [SW92]. The name of the game is to compute a basis H such that $w(H) = w(\Gamma)$, and this can be done efficiently when the combinatorial dimension is bounded [Cla95, MSW96, Sei91]. Lemma 8.7 has an analogue for LP-type problems which extends into a two-ways relation between Helly numbers and the combinatorial dimension of LP-type problems [Ame94].

9 HELLY NUMBERS OF ACYCLIC FAMILIES

In this chapter, we prove a new condition under which Helly numbers can be bounded. Call a family $\mathcal{F}$ of open subsets of $\mathbb{R}^d$ *acyclic* if for any non-empty sub-family $G \subseteq \mathcal{F}$, each connected component of the intersection of the elements of G is a $\mathbb{Q}$-homology cell. We show the following:

Theorem 9.1. *Let $\mathcal{F}$ be a finite acyclic family of open subsets of $\mathbb{R}^d$. If any sub-family of $\mathcal{F}$ intersects in at most r connected components then the Helly number of $\mathcal{F}$ is at most $r(d+1)$.*

Our proof uses a new object, the *multinerve* $\mathcal{M}(\mathcal{F})$ of $\mathcal{F}$, a simplicial poset (see below) that encodes the intersection pattern of $\mathcal{F}$ more finely than the nerve. Considering multinerves, we avoid the need for an intersectional structure such as the (G, r)-families of the Grünbaum-Motzkin setting.

We establish Theorem 9.1 by following the approach outlined in Chapter 8. We first associate to every simplicial poset X some index $J(X)$, closely related to the Leray number, which we control via an analogue of the Nerve theorem for multinerves of acyclic families. We then project the multinerve of $\mathcal{F}$ onto its nerve, and extend Theorem 8.5 to bound the Leray number of the nerve in terms of the J-index of the multinerve and the multiplicity of the projection.

Theorem 9.1 does not readily bound the transversal Helly number of families of disjoint unit balls. We do, however, prove a more general statement (Theorem 9.7) that sharpens our previous upper bound on $\mathcal{H}_d$ from $4d-1$ to $4d-2$ when $d \geqslant 6$.

We present the main steps of the proof of Theorem 9.1 but refer to [19] for the, rather technical, details of the proofs.

⋆ ⋆ ⋆

Intuitively, a simplicial partially ordered set (simplicial poset for short) is a set of simplices with an incidence relation; a d-simplex still has $d+1$ distinct vertices; however, in contrast to simplicial complexes, there may be several simplices with the same vertex set, but no two can be incident to the same higher-dimensional simplex.

Formally, let X be a finite set and $\preceq$ a partial order on X; we also say that $(X, \preceq)$ is a partially ordered set, or that X is a *poset* to save breath. Let $[\alpha, \beta] = \{\tau \in X \mid \alpha \preceq \tau \preceq \beta\}$ denote the *segment* defined by α and β. A map $\varphi : X \to Y$ between two posets $(X, \preceq^X)$ and $(Y, \preceq^Y)$ is *monotone* if it preserves the order: for any $\sigma, \tau \in X$ $\sigma \preceq^X \tau \Rightarrow \varphi(\sigma) \preceq^Y \varphi(\tau)$. An *isomorphism* of posets is a monotone bijection between them.

A poset X is a *simplicial poset* if it satisfies two conditions. First, X must have a *least element* 0, that is $0 \preceq \sigma$ for any $\sigma \in X$. Second, for any $\sigma \in X$, there must exist some integer d such that the lower segment $[0, \sigma]$ is isomorphic to $2^{[d]}$, the poset of faces of a d-simplex partially ordered by the inclusion; d is then called the *dimension* of σ. The elements of X are called its *simplices* and the simplices of dimension 0 (i.e. that only dominate 0) are its *vertices*. If $\tau \preceq \sigma$ we also say that τ is *contained in* (or *a face of*) σ.

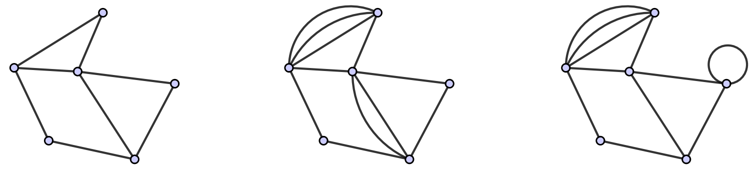

Figure 48: Simplicial complex (left) vs simplicial poset (middle). Simplicial sets (right) are even more general structures.

Simplicial posets lie in between simplicial complexes and the more general notion of *simplicial sets* used in algebraic topology [May92, GJ99]. The simplices of a simplicial complex, ordered by inclusion, form a simplicial poset (with $\emptyset$ as least element). The converse is not always true: the one-dimensional simplicial complexes are precisely the graphs without loops or multiple edges, while the one-dimensional simplicial posets correspond to the graphs without loops but possibly with multiple edges (see Figure 48). Let τ be a simplex of a simplicial poset with set of vertices V. The map that associates to any face of τ the set of vertices of that face is a bijection between $[0, \tau]$ and 2^V. There may, however, exist several simplices with the same set of vertices, but no two of them can be faces of one and the same simplex.

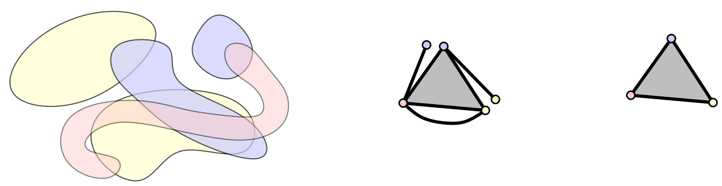

Figure 49: Left: A family $\mathcal{F}$ of subsets of $\mathbb{R}^2$. Middle: Its multinerve $\mathcal{M}(\mathcal{F})$. Right: Its nerve $N(\mathcal{F})$.

Let $\mathcal{F}$ be a finite family of subsets of a topological space. We define the *multinerve* $\mathcal{M}(\mathcal{F})$ of $\mathcal{F}$ as:

$$\mathcal{M}(\mathcal{F}) = \left\{ (C, A) \;\middle|\; A \subseteq \mathcal{F}, C \text{ is a connected component of } \bigcap_A \right\}.$$

By convention, we put $\bigcap_\emptyset = \bigcup_\mathcal{F}$, and in particular, $(\bigcup_\mathcal{F}, \emptyset)$ belongs to $\mathcal{M}(\mathcal{F})$. We turn $\mathcal{M}(\mathcal{F})$ into a poset by equipping it with the partial order:

$$(C', A') \preceq (C, A) \Leftrightarrow C \supseteq C' \text{ and } A \subseteq A'.$$

To get an intuition, it does not harm to assume that, whenever A and A' are different subsets of $\mathcal{F}$, the connected components of $\bigcap_A$ and of $\bigcap_{A'}$ are different. Under this assumption, $\mathcal{M}(\mathcal{F})$ can be identified with the set of all connected components of the intersection of any sub-family of $\mathcal{F}$, equipped with the opposite of the inclusion order. See Figure 49 for an example.

Lemma 9.2. *The poset $\mathcal{M}(\mathcal{F})$ is simplicial.*

Proof. The projection on the second coordinate identifies any lower segment $[(\bigcup_\mathcal{F}, \emptyset), (C, A)]$ with the simplex 2^A. Indeed, let $A' \subseteq A$ and let $C' \subseteq \bigcup_\mathcal{F}$. The lower segment $[(\bigcup_\mathcal{F}, \emptyset), (C, A)]$ contains (C', A') if and only if C' is the connected component of $\bigcap_{A'}$ containing C. Moreover, by definition, $\mathcal{M}(\mathcal{F})$ contains a least element, namely $(\bigcup_\mathcal{F}, \emptyset)$. The statement follows. □

Intuitively, $\mathcal{M}(\mathcal{F})$ is an "expanded" version of $N(\mathcal{F})$: while $N(\mathcal{F})$ has one simplex for each non-empty intersecting sub-family, $\mathcal{M}(\mathcal{F})$ has one simplex for each *connected component* of an intersecting sub-family.

⋆ ⋆ ⋆

Our proof, like the one by Kalai and Meshulam, controls Helly numbers through a homological index. The *homology* of a simplicial poset can be defined in three different ways: as a direct extension of simplicial homology for simplicial complexes, as a special case of simplicial homology of simplicial sets [GJ99, May92], or via the singular homology of its *geometric realization*; all three definitions are equivalent in that they lead to isomorphic homology groups. We refer to Appendix B for more details on these notions. If X is a simplicial poset with vertex set V and $S \subseteq V$, the *induced simplicial sub-poset* $X[S]$ is the poset of elements of X whose vertices are in S, ordered by the order of X. Along with the notion of induced simplicial sub-poset and homology groups, the notion of *Leray number* extends immediately to simplicial posets:

$$L(X) = \min\{i \in \mathbb{N} \mid \forall S \subseteq V, \forall j \geqslant i, \quad \tilde{H}_j(X[S]) = 0\}.$$

We bound $L(\mathcal{M}(\mathcal{F}))$ from above via the following analogue of the Nerve Theorem for acyclic families.

Theorem 9.3 ([19, Section 3])**.** *If $\mathcal{F}$ is an acyclic family in $\mathbb{R}^d$ then $\tilde{H}_\ell(\mathcal{M}(\mathcal{F})) \cong \tilde{H}_\ell(\bigcup_\mathcal{F})$ for any $\ell \geqslant 0$.*

Proof principle. A classical theorem of Leray [Bre97, God73, KS90, Spa66] states that the Čech complex of an acyclic cover captures the homology of its union. This theorem follows from Leray's theorem via an interpretation of a multinerve, or more precisely its chain complex, as a Čech complex. □

If $\mathcal{F}$ is an acyclic family of open sets in $\mathbb{R}^d$ then $\bigcup_\mathcal{F}$ has trivial homology in dimension d or larger [Gre67, p. 121] and the same therefore holds for $\mathcal{M}(\mathcal{F})$. Since induced subposets of $\mathcal{M}(\mathcal{F})$ are themselves multinerves of an acyclic family, it follows that the Leray number of $\mathcal{M}(\mathcal{F})$ is at most d.

Our projection theorem considers the following refinement of the Leray number. Given a simplicial poset X we let $J(X)$ be the smallest integer ℓ such that for every $j \geqslant \ell$, every $S \subseteq V$, and every simplex σ of $X[S]$, we have $\tilde{H}_j(\dot{D}_{X[S]}(\sigma)) = 0$. Here $\dot{D}_Y(\sigma)$ denotes the *order complex* of $(\sigma, \cdot]$ in Y, that is the simplicial complex consisting of all finite chains of $(\sigma, \cdot]$, ordered by inclusion. It is simple, if a bit technical, to show that the index $J(\mathcal{M}(\mathcal{F}))$ of an acyclic family $\mathcal{F}$ of open sets in $\mathbb{R}^d$ is also at most d [19, Lemma 19].

Let $\phi : X \to Y$ be a monotone map between two simplicial posets. We say that ϕ is *dimension-preserving* if for any $\sigma \in X$ the dimension of $\varphi(\sigma)$ equals the dimension of σ; the *multiplicity* of ϕ is the maximum number of elements in a fiber, that is $\max_{\tau \in Y} |\phi^{-1}(\tau)|$. When X and Y are simplicial complexes, these definitions coincide with those given in Chapter 8. Our projection theorem is the following.

Theorem 9.4 ([19, Theorem 15])**.** *If $\phi : X \to Y$ is a surjective, dimension-preserving monotone map of multiplicity r from a simplicial poset X to a simplicial complex Y then $J(Y) + 1 \leqslant r(J(X) + 1)$.*

Proof principle. The special case where X is a simplicial complex was proven by Kalai and Meshulam [KM08, Theorem 1.3] in a slightly different terminology. Their proof relies on a characterization of Leray numbers in terms of the homology of links [KM06, Proposition 3.1]; the main technical difficulty in extending their proof to the simplicial poset setting is that it is not clear whether there is a generalization of the notion of links to simplicial posets that leads to a similar characterization. We introduced the index J as a way around that technical difficulty; the rest of the proof extends, mutatis mutandis. □

We note that already in the special case where X is a simplicial complex the bound on $L(Y)$ is tight (see the remark after Theorem 1.3 of [KM08]).

We can now complete the proof of Theorem 9.1. Let $\mathcal{F}$ be a finite acyclic family of open subsets of $\mathbb{R}^d$. As explained above, Theorem 9.3 implies that $J(\mathcal{M}(\mathcal{F}))$ is at most d. The projection π on the second coordinate, defined by $\pi((C, A)) = A$, is a monotone map from $\mathcal{M}(\mathcal{F})$ to $N(\mathcal{F})$. Since $\dim(C, A) = |A| - 1$ we also have that π is dimension preserving. Also, every fiber $\pi^{-1}(A)$ consists of those (C, A) where C is a connected component of $\bigcap_A$. If every sub-family of $\mathcal{F}$ intersects in at most r connected components then π has multiplicity at most r and Theorem 9.4 yields:

$$L(N(\mathcal{F})) \leqslant rJ(\mathcal{M}(\mathcal{F})) + r - 1 \leqslant rd + r - 1.$$

It follows that the Helly number of $\mathcal{F}$ is at most $r(d + 1)$.

⋆ ⋆ ⋆

Let us now turn our attention to the transversal Helly numbers $\mathcal{H}_d$. Let $\mathcal{F}$ be a family of disjoint unit *open* balls in $\mathbb{R}^d$ and let $\mathcal{F}^{\mathcal{T}}$ denote the family of non-oriented line transversals to the members of $\mathcal{F}$. Applying Theorem 9.1 to the family $\mathcal{F}^{\mathcal{T}}$ requires to address a few technicalities that lead to generalize it into Theorem 9.7.

The first issue to address is that the elements of $\mathcal{F}^{\mathcal{T}}$ are not, naturally, subsets of some $\mathbb{R}^m$ but of the Grassmannian $G_{2,d+1}$. The nature of the ambient space is used in two places in the proof of Theorem 9.1. First, it ensures, in the proof of Theorem 9.3, that the connected components and the arc-wise connected components of any open subset of the space agree; the same holds, in fact, in any locally arc-wise connected topological space. Second, it ensures that $\bigcup_{\mathcal{F}}$, and therefore $\mathcal{M}(\mathcal{F})$, has trivial homology in dimension d and higher. We can easily extend this to arbitrary topological spaces Γ by introducing an index d_Γ defined as the smallest integer such that every open subset of Γ has trivial $\mathbb{Q}$-homology in dimension d_Γ and higher; the bound on the Helly number then becomes $r(d_\Gamma + 1)$.

The next question is to determine $d_{G_{2,d+1}}$. If Γ is a k-dimensional manifold the index d_Γ is known [Gre67, p. 121]: it is k if Γ is non-compact or non-orientable and $k + 1$ otherwise. Since $G_{2,d+1}$ is a compact, orientable manifold of dimension $2d - 2$ we therefore have $d_{G_{2,d+1}} = 2d - 1$. We can, in fact, work in the space E of lines intersecting some open ball B in $\mathbb{R}^d$ that contains all balls in $\mathcal{F}$. Since E is a non-compact $2d - 2$-manifold, its associated index is $d_E = 2d - 2$, one less than $d_{G_{2,d+1}}$.

The second issue to address is that $\mathcal{F}^{\mathcal{T}}$ is not quite acyclic. For any subfamily $G \subseteq \mathcal{F}$ the homology of $\bigcap_{G^{\mathcal{T}}}$ can be described using the Convexity theorem of Chapter 2 and Lemma 3.1: if $|G| = 1$ then $\bigcap_{G^{\mathcal{T}}}$ has the same homotopy type as $\mathbb{RP}^{d-1}$, else it has up to k connected components, each contractible, where k is the maximum number of geometric permutations of G. The family $\mathcal{F}^{\mathcal{T}}$ therefore fails to be acyclic, but only because the members have non-trivial homology; when we consider intersection of 2 or more elements that homology disappears. Call a finite family C of subsets of a topological space *acyclic with slack* s if for every subfamily $G \subseteq C$ and every $i \geqslant \max(1, s-|G|)$ we have $\tilde{H}_i(\bigcap_G, \mathbb{Q}) = 0$. We can extend theorem 9.3 as follows:

Theorem 9.5 ([19, Theorem 8]). *Let* C *be a family of open sets in a locally arc-wise connected topological space* Γ*. If* C *is* acyclic with slack s *then* $\tilde{H}_\ell(\mathcal{M}(C)) \cong \tilde{H}_\ell(\bigcup_C)$ *for any non-negative integer* $\ell \geqslant s$*.*

We only use the Multinerve theorem to control the Leray number of the multinerve (or its variant, the index J), which only cares about homology in *high* dimension. We can thus generalize Theorem 9.1 to families that are acyclic with slack s, the bound on the Helly number becoming $r(\max(d_\Gamma, s)+1)$. Observe that Chapter 2 and Lemma 3.1 now imply that $\mathcal{F}^{\mathcal{T}}$ is acyclic with slack $d+1$.

To apply our current generalization of Theorem 9.1 it remains to bound r, which is the maximum number of geometric permutations of a subfamily of $\mathcal{F}$. This number is at most 3 by Theorem 4.1. Altogether, we obtain the bound

$$\mathcal{H}_d \leqslant 3(\max(2d-2, d+1)+1) = 6d-3 \quad \text{for } d \geqslant 3,$$

which is not very exciting in the light of the bound $\mathcal{H}_d \leqslant 4d-1$ obtained in Part i.

Theorem 4.1 states that a subfamily of $\mathcal{F}$ of *at least* 9 balls has at most 2 geometric permutations. We can take advantage of this by allowing, in Theorem 9.1, the intersections of small subfamilies to have more than r connected components. We do this by proving another projection theorem:

Lemma 9.6 ([19, Lemma 20]). *Let* X *and* Y *be two simplicial posets and* $k \geqslant 0$*. Assume that there exists a monotone, dimension-preserving and surjective map* $f : X \to Y$ *between two simplicial posets and that the restriction of* f *to the simplices of* X *of dimension at least* k *is a bijection onto the simplices of* Y *of dimension at least* k*. Then* $J(Y) \leqslant \max\big(J(X), k+1\big)$.

Now, let $k \geqslant 2$ be some integer. We construct a simplicial poset $\mathcal{M}_{red}(\mathcal{F}^{\mathcal{T}})$ by identifying together two simplices of $\mathcal{M}(\mathcal{F}^{\mathcal{T}})$ if and only if they are of the form (C, A) and (C', A') with $A = A'$ and A has dimension at most $k-2$. The natural projection

$\mathcal{M}(\mathcal{F}^{\mathcal{T}}) \to \mathcal{M}_{red}(\mathcal{F}^{\mathcal{T}})$ satisfies the conditions of Lemma 9.6, and $J(\mathcal{M}_{red}(\mathcal{F}))$ is at most $\max(d, k)$. We finally obtain the following generalization of Theorem 9.1:

Theorem 9.7. *Let $\mathcal{F}$ be a finite family of open subsets of a locally arc-wise connected topological space Γ. If (i) $\mathcal{F}$ is acyclic with slack s and (ii) any sub-family of $\mathcal{F}$ of cardinality at least t intersects in at most r connected components then the Helly number of $\mathcal{F}$ is at most $r(\max(d_\Gamma, s, t) + 1)$.*

We can apply this theorem to $\mathcal{F}^{\mathcal{T}}$ with $r = 2$ by setting $d_\Gamma = 2d-2$, $s = d+1$ and $t = 9$. For $d \leqslant 5$ the bound is not interesting, but for $d \geqslant 6$ we obtain:

$$\mathcal{H}_d \leqslant 4d - 2 \quad \text{for } d \geqslant 6 \tag{9.1}$$

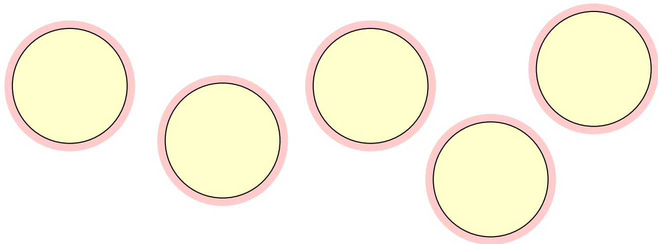

Figure 50: An open thickening of a family of disjoint closed unit balls.

One could object that Equation (9.1) holds for the transversal Helly number of disjoint open unit balls, whereas what we studied in Part i was the transversal Helly number of disjoint closed unit balls. The gap can be bridged by considering thickenings. An *open thickening* of a subset H of $\mathbb{R}^d$ is a family (H^ε) such that (i) any H^ε is an open set, (ii) if $\varepsilon < \varepsilon'$, then $H^\varepsilon \subseteq H^{\varepsilon'}$, and (iii) $\bigcap_{\varepsilon>0} H^\varepsilon = H$. For a family G of subsets of $\mathbb{R}^d$, we let $G^\varepsilon = \{H^\varepsilon \mid H \in G\}$.

Lemma 9.8. *Let H be a finite family of compact convex sets in $\mathbb{R}^d$ and H^ε be an open thickening of H. There exists $\varepsilon > 0$ such that for every $G \subseteq H$, the family G has a line transversal if and only if the family G^ε has a line transversal.*

Proof. Let $G \subseteq H$. We argue that G has a line transversal if and only if G^ε has a line transversal for arbitrarily small ε. This implies the statement as H is finite.

That the condition is necessary is obvious. Conversely, assume that there exists a sequence (ε_n) decreasing towards zero, and, for every n, a line (ℓ_n) transversal to G^{ε_n}. For each $A \in G$ we pick a point $p_{A,n}$ in $A^{\varepsilon_n} \cap \ell_n$. Up to taking a subsequence, we can assume that (ℓ_n) converges towards a line ℓ, and that each sequence $(p_{A,n})$ converges towards some point p_A (by compactness of $G_{2,d+1}$, and since the objects are bounded). Of course, each p_A belongs to ℓ, and also to the closure of each A^{ε_n}, hence to A, since A is closed. So ℓ is a line transversal to G. □

Now, any family $\mathcal{F}$ of disjoint closed unit balls in $\mathbb{R}^d$ admits an open thickening $(\mathcal{F}^\varepsilon)$ by families of disjoint open equal-radius balls; Lemma 9.8 that the transversal Helly number of $\mathcal{F}$ is equal to the transversal Helly number of one of the $\mathcal{F}^\varepsilon$. Equation (9.1) therefore extends from families of disjoint open unit balls to families of disjoint closed unit balls.

⋆ ⋆ ⋆

NOTES.

This chapter is based on an article co-authored with Éric Colin de Verdière and Grégory Ginot [19].

Families of disjoint translates of a planar convex figure, like families of disjoint unit balls, have convex cones of directions and a bounded number of geometric permutation [Asi99, KLL87]. We can thus use Theorem 9.7 to bounds the transversal Helly number of these families; the best choice of parameters (s, t) gives a bound of 10. This may seem bad as the sharp bound is 5, as proven by Tverberg [Tve89], as conjectured by Grünbaum in 1957. Let us point out, however, that the first bound, obtained by Katchalski [Kat86], was 128 and was only published three decades later.

The closest predecessor of the statement of Theorem 9.1 is a theorem of Matoušek [Mat97] that bounds by some function $h(d, r)$ the Helly number of any family of sets in $\mathbb{R}^d$ such that the intersection of any sub-family has at most r connected components, each of which is $(\lceil d/2 \rceil - 1)$-connected.[1] His proof, however, only gives a loose bound on the Helly number (in fact, no explicit bound is given), whereas our approach gives sharp, explicit, bounds. Interestingly, his theorem allows the connected components to have nontrivial *homotopy* in *high* dimension, whereas Theorem 9.7 lets them have nontrivial *homology* in *low* dimension. A natural question is whether these two types of assumption could be unified in some way.

1 A set is k-connected if its ith homotopy group vanishes for $i \leqslant k$.

Part IV

PERSPECTIVES

PERSPECTIVES

C'est pas un peu fatigant, à la longue, de réfléchir
comme ça tout le temps ?
H. P.

We conclude this promenade by a few thoughts on some new questions raised by the results we presented.

⋆ ⋆ ⋆

In the light of our results, and in particular of the Convexity theorem (Theorem 2.1 of Chapter 2), families of disjoint Euclidean balls appear peculiar when it comes to the geometry of line transversals. This impression should be tempered by the fact that we only use this convexity structure as a shortcut to weaker properties that may perhaps generalize beyond disjoint balls.

Several of our arguments rely on the fact that connected components of line transversals to 2 or more disjoint balls are contractible (Lemma 3.1). Whether this extends to families of disjoint convex sets is unclear; we certainly know of no example showing otherwise.

Question 9.1. *Does there exist a family of $n \geqslant 2$ disjoint convex sets in $\mathbb{R}^3$ with a connected component of line transversals that is not contractible?*

While this question is easy for $n = 2$ (no such pair exists), it seems open already for $n = 3$. If the answer turns out to be negative, our general Helly-type theorems (Theorems 9.1 and 9.7) would reduce the search for bounded transversal Helly numbers to a search for conditions ensuring that the number of connected components of line transversals remain bounded.

The main consequence of the contractibility of connected components of line transversals is the pinning theorem of Chapter 3 (Theorem 3.5). We could prove upper bounds on the pinning numbers in situations where the Helly number is known to be unbounded: disjoint balls of arbitrary radii (Theorem 3.5), intersecting balls in $\mathbb{R}^3$ (Chapter 7), polytopes in sufficiently generic position in $\mathbb{R}^3$ (Chapter 5). So far, we found no situation where more than six disjoint convex sets in $\mathbb{R}^3$ are needed to immobilize a line. This is in sharp contrast with the Helly numbers of sets of line transversals, which are only bounded in very specific settings (disjoint translates of a convex set in the plane, disjoint balls of bounded radius disparity in $\mathbb{R}^d$).

Question 9.2. *Does there exist a family of $n \geqslant 7$ disjoint convex sets in $\mathbb{R}^3$ that minimally pin a line?*

Answering this question in the case of disjoint convex polytopes would be a natural starting point. The main difficulty is that some non-convexity hides in situations where the immobilized line passes through a face of a polytope, as the set of line intersecting that polytope in the vicinity of the line is isometric to (the intersection of the Klein quadric with) the union of two halfspaces. The case of general convex sets is likely to be challenging; already settling the conjecture in the case of semi-algebraic convex sets would be interesting.

Another natural question (already mentioned p. 81) is whether differential geometry considerations such as those used to study pinning theorems in Chapter 7 could yield new insight on other questions in line geometry. A natural candidate problem is the characterization of quadruples of objects in $\mathbb{R}^3$ with infinitely many common tangent lines. When the objects are balls or lines, this question was answered by closely inspecting algebraic systems describing these common tangents [6, MPT01, Meg01, The02, MST03, MS05, ST08]; this approach, however, does not extend easily even to the case of four ellipsoids in $\mathbb{R}^3$.

Question 9.3. *What configurations of four smooth rotund convex sets in $\mathbb{R}^3$ have infinitely many common tangent lines?*

A necessary condition for a quadruple of smooth rotund convex sets to have infinitely many common tangents is that they are in non-transverse position with respect to all of their common tangents, considerably narrowing down the range of configurations to analyze.

⋆ ⋆ ⋆

Hadwiger's transversal theorem is a relative of Danzer's theorem that asserts that a family $\mathcal{F}$ of disjoint convex sets has a line transversal if any three sets can be met by a line *consistently with some ordering on* $\mathcal{F}$. Every simplex in the multinerve $\mathcal{M}(\mathcal{F}^{\mathcal{T}})$ of $\mathcal{F}^{\mathcal{T}}$ is associated with a connected component of line transversal to some subfamily of $\mathcal{F}$. If we fix some ordering $\prec$ on $\mathcal{F}$, the simplices of $\mathcal{M}(\mathcal{F}^{\mathcal{T}})$ whose associated transversals are consistent with $\prec$ form not only a simplicial poset, but a simplicial complex. This simplicial complex is, unfortunately, neither a nerve nor an induced sub-poset. It is thus not clear whether its Helly number can be bounded from the Leray number of $\mathcal{M}(\mathcal{F}^{\mathcal{T}})$; this would be interesting as such a bound would bring, for the first time, transversal Helly numbers and Hadwiger-type theorems under the same umbrella.

⋆ ⋆ ⋆

Helly numbers are related, as mentioned in the notes of Chapter 8, to the notion of *combinatorial dimension* of *LP-type* problems. Pinning numbers, as local analogues of Helly numbers, enjoy a similar interpretation. Specifically, call a cylinder C a *locally* smallest enclosing cylinder of a point set P if any enclosing cylinder of P with direction sufficiently close to that of C must have strictly greater radius; Theorem 3.5 implies that if a locally smallest enclosing cylinder of a set of points $P \subset \mathbb{R}^d$ has radius at most half the smallest inter-point distance then there must be some at most $2d-1$ points in P for which C is already a locally smallest enclosing cylinder. In the light of the robustness of pinning theorems, the sparsity condition (relating the radius of C to the smallest inter-point distance in P) may be unnecessary.

When the combinatorial dimension of a LP-type problem can be bounded then this problem can be solved in time linear in its number of constraints. For example, our positive answer to Danzer's conjecture implies that in any fixed dimension, the smallest enclosing cylinder to a *sparse point set*[2] can be computed in $O(n)$ time; this is in sharp contrast with the general case, where already in $\mathbb{R}^3$ the best known algorithm takes $O(n^{3+\epsilon})$-time [AAS99]. Can a similar algorithmic framework exploit bounded pinning numbers?

2 A point set is sparse if the radius of its smallest enclosing cylinder is at most half the smallest inter-point distance.

Part V

APPENDICES

A

LINE SPACE REPRESENTATION

While the geometry of lines has been studied for quite some time, as witnessed by Euclid's *elements*, the modern view of line geometry, studying the space of line as an object in itself, is more recent; it finds its origin in the second half of the 19^{th} century, with the work of people such as Plücker, Klein and Grassmann.

Let us denote by G the space of lines in $\mathbb{R}^3$. We are interested in *representing* G, that is, defining a coordinate system that preserves the topology of G; in other words, we want to find some homeomorphism between G and some parameter space. But what topology should we equip G with? The answer is not unique, but a natural choice comes from the identification

$$G \simeq \left(\mathbb{R}^3 \times \mathbb{R}^3 \setminus \Delta\right) / \sim,$$

where $\Delta = \{(p,p) \mid p \in \mathbb{R}^3\}$ and $(a,b) \sim (c,d)$ if and only if a, b, c and d are on the same line. The standard topology on $\mathbb{R}^3 \times \mathbb{R}^3 \setminus \Delta$ indeed defines, via the quotient by $\sim$, the topology on G where the open subsets are the sets of lines transversals to pairs of open subsets of $\mathbb{R}^3$.

Representing G requires some care, and various "natural" choices of parameters introduce singularities in the representation. One may, for instance, pick two planes, together with coordinate systems on these planes, and map a line to the quadruple obtained by juxtaposing the coordinate of its intersection points with these two planes (c.f. page 53). This parameterization may be very useful (as demonstrated in Chapter 5) but only works if one only cares about what happens near a line intersecting the planes in two distinct points; lines parallel to one of the planes or passing through the intersection of the two planes are, however, not represented. An alternative could be to fix, for each direction $u \in \mathbb{S}^2$, some coordinate system (O, x_u, y_u) for the plane Π_u with normal u and passing through the origin; we can then associate to each line ℓ two coordinates representing its direction vector $\overrightarrow{\ell}$ and the two coordinates of $\ell \cap \Pi_u$. Such a system may sometimes be convenient [Duroo] but does, nevertheless, present singularities (both in the representation of the direction using two coordinates, and in the discontinuity that must exist in the function $u \mapsto (x_u, y_u)$ by the hairy ball theorem).

The above attempts at representing the space G of lines in $\mathbb{R}^3$ show that G is a 4-dimensional manifold. This manifold is not, however, homeomorphic to $\mathbb{R}^4$, so no representation by $\mathbb{R}^4$ is possible. One may, however, map the line $\ell = p + \mathbb{R}u$, where p is a point and u a vector, to the sextuple $\xi(\ell) = (u; \overrightarrow{Op} \times u)$, where $(a; b)$ denotes the juxtaposition of the coordinates of two vectors a and b. Note that if $p + \mathbb{R}u$ and $p' + \mathbb{R}u'$ are the same line then the vectors

$$(u; \overrightarrow{Op} \times u) \quad \text{and} \quad (u'; \overrightarrow{Op'} \times u')$$

are equal up to multiplication by a constant. This suggests to consider $\xi(\ell)$ not as a point in $\mathbb{R}^6$, but as a point in $\mathbb{RP}^5$. Observe that if $\xi(\ell) = (a; b)$ then $a \cdot b = 0$ and a is not zero; in other words, ξ maps every line in $\mathbb{R}^3$ to a point $(x_1, \ldots, x_6) \in \mathbb{RP}^5$ with

$$x_1x_4 + x_2x_5 + x_2x_6 = 0 \quad \text{and} \quad (x_1, x_2, x_3) \neq (0, 0, 0). \quad \text{(A.1)}$$

It is not hard to check that the converse is true: any point of $\mathbb{RP}^5$ satisfying Condition (A.1) is the image through ξ of some line in $\mathbb{R}^3$. This representation is, in a nutshell, the affine rephrasing of the standard coordinate system introduced by Julius Plücker, and thus known as *Plücker coordinates*, and their interpretation, due to Felix Klein, as a map between lines in $\mathbb{RP}^3$ and a hyperquadric in $\mathbb{RP}^5$ (often referred to as the *Klein quadric*). Our parameterization of lines using points in $\mathbb{R}^5$ (page 53) is simply the image, under some affine transformation, of (a subset of) the Plücker-Klein model.

The Klein quadric (Γ) can also be written as $\Gamma = \{q(u, u) = 0 \mid u \in \mathbb{RP}^5\}$, where q is the bilinear form

$$q(u, v) = u_1v_4 + u_2v_5 + u_3v_6 + u_4v_1 + u_5v_2 + u_6v_3.$$

Associated to any bilinear form is a notion of orthogonality: u is orthogonal to v with respect to q if $q(u, v) = 0$. In the Plücker-Klein model, the orthogonality with respect to q has a simple geometric interpretation: two lines in $\mathbb{RP}^3$ meet if and only if their Plücker coordinates are orthogonal with respect to q. Since q is bilinear, the set of points orthogonal to a fixed point u recasts as a hyperplane; in our affine model, the hyperplane bounding $\bar{\mathcal{U}}_g$ is simply the set of points orthogonal to the point representing the constraint g.

Plücker coordinates were later generalized by Grassmann to represent the space of k-dimensional flats in $\mathbb{RP}^d$, leading to the development of exterior calculus. Already for lines in $\mathbb{RP}^d$, however, things become more complicated as the Grassmann coordinates no longer live on a hypersurface but on a lower-dimensional manifold. We refer to [PW01, Chapter 2] for a more detailed treatment of the Plücker-Klein model, and to [HP94] for the more general theory of Grassmannian manifolds.

B
A PRIMER ON HOMOLOGY THEORY

This appendix presents a brief introduction to homology theory. Our goal here is merely to equip the reader not already acquainted with this branch of mathematics with the definitions and basic intuition of the objects manipulated in Chapters 8 and 9. For a more thorough exposition of this topic we refer to one of the classic textbooks, e.g. [Hat02].

We first define *simplicial* homology of simplicial complexes. We then introduce the *singular homology* of the *geometric realizations* of simplicial complexes, and extend all these notions to simplicial posets. Recall that *simplicial complex* and *simplicial posets* were defined page 85 and 97, respectively. Let us emphasize that we only consider homology with *rational* coefficients as that is what we use in Chapters 8 and 9.

⋆ ⋆ ⋆

Let X be a simplicial complex and assume chosen an ordering on its set V of vertices: $V = \{v_1, \dots, v_\ell\}$. For $n \geqslant 0$, let $C_n(X)$ be the $\mathbb{Q}$-vector space with basis the set of simplices of X of dimension n; in particular note that $(C_n(X), +)$ is a group. An element of $C_n(X)$ is called a *simplicial n-chain*, a n-chain for short.

Let $n \geqslant 0$. For $i = 0, \dots, n$ we define the i^{th} *face* of a n-simplex σ as the face $d_i(\sigma)$ obtained by removing from σ its i^{th} vertex. For instance:

$$d_0(\{v_2, v_5\}) = \{v_5\} \quad \text{and} \quad d_1(\{v_2, v_5\}) = \{v_2\}.$$

The map d_i extends, by linearity, into a map from $C_n(X)$ to $C_{n-1}(X)$, called the i^{th} *face operator*; for instance:

$$d_0\left(\{v_2, v_5\} + \frac{4}{7}\{v_5, v_7\}\right) = \{v_5\} + \frac{4}{7}\{v_7\}.$$

We then define the *boundary operator* as the map $d = \sum_{i=0}^{n}(-1)^i d_i$ from $C_n(X)$ to $C_{n-1}(X)$:

$$d\left(\{v_2, v_5\} + \frac{4}{7}\{v_5, v_7\}\right) = \frac{3}{7}\{v_5\} + \frac{4}{7}\{v_7\} - v_2.$$

There are now two types of simplicial chains in $C_n(X)$ of particular interest. Let

$$B_n(X) = \{d(s) \mid s \in C_{n+1}(X)\}$$

denote the set of *boundaries*, i.e. those chains that are the image under d of some $(n+1)$-chain. Also, let

$$Z_n(X) = \{s \mid s \in C_n(X) \text{ and } d(s) = 0\}$$

denote the set of *cycles*, i.e. those chains whose images under d is zero. The linearity of d implies that both B_n and Z_n are groups (for the addition).

Consider the map $d^2 : C_n(X) \to C_{n-2}(X)$ obtained by composing d with itself. Let σ be a n-simplex and let v and v' be its i^{th} and j^{th} vertices, with $i < j$. If we expand $d^2(\sigma)$ we obtain $\sigma \setminus \{v, v'\}$ twice: once with sign $(-1)^{i+j}$, for the removal of v' then v, and once with sign $(-1)^{i+j-1}$ for the removal of v then v'. These two contributions therefore cancel each other out, and the resulting coefficient of $\sigma \setminus \{v_i, v_j\}$ in $d^2(\sigma)$ is zero. This extends by linearity into $d^2 = 0$ over all of $C_n(X)$, and it follows that any boundary is a cycle, that is $B_n \subseteq Z_n$. The converse is often not true, and this is precisely what the homology groups capture: which cycles are *not* boundaries. Formally, the n*-th homology group* of X is the quotient $H_n(X) = Z_n(X)/B_n(X)$.

Example B.1. Let X be a simplicial complex and let S denote its 1-skeleton, that is the subset of X of simplices of dimension at most 1. If v and v' are two vertices in the same connected component of G (seen as a graph) then there exists an oriented path $(v^0, v^1), (v^1, v^2, \dots, (v^{k-1}, v^k)$ in G from $v^0 = v$ to $v^k = v'$. Letting $\varepsilon_i = 1$ if v^i precedes v^{i+1} in the ordering on V, and $\varepsilon_i = -1$ otherwise, we have

$$d\left(\sum_{0 \leqslant i \leqslant k} \varepsilon_i \{v^i, v^{i+1}\}\right) = v^k - v^0 = v' - v.$$

By linearity, any 0-chain of X consisting of two vertices of the same connected component with opposite weights is a boundary. In fact, a 0-chain is a boundary if and only if for every connected component γ of S, the coefficients (in the chain) of the vertices of γ sum to 0. It follows that $B_0(X) \simeq \mathbb{Q}^{\ell - c}$ where ℓ is the number of vertices in X and c the number of connected components of S. On the other hand, d maps any vertex of X to 0, and thus $Z_0(X) \simeq \mathbb{Q}^{\ell}$. It follows that $H_0(X) = Z_0(X)/B_0(X) \simeq \mathbb{Q}^c$ "counts" the number of connected components of (the 1-skeleton of) X.

Example B.2. Let $X = 2^V$ be the k-dimensional simplex and let $Y = X \setminus \{V\}$ be a simplicial hole with $k+1$ vertices. The chain

$$s = d(V) = \sum_{0 \leqslant i \leqslant k+1} (-1)^i V \setminus \{v_i\}$$

is a boundary in X and therefore a cycle in X. In Y, s remains a cycle (as it uses only simplices from Y) but is no longer a boundary (as we removed precisely the simplex it was the boundary of). It follows that $B_{k-1}(Y) \subsetneq Z_{k-1}(Y)$, and $H_{k-1}(Y)$ is non-trivial. In other words, a simplicial hole of cardinality $k+1$ has non-trivial homology in dimension $k-1$. For $n \geqslant k$, since Y has no simplex of dimension n, $C_n(Y)$ is empty and $H_n(Y)$ is trivial.

⋆ ⋆ ⋆

The topological meaning of a simplicial complex having non-trivial homology groups may be more apparent if one considers the *singular* homology of their geometric realization. There are two classical ways to define such "geometric realization".

The first approach is via geometric simplicial complex. A *geometric n-simplex* is the convex hull of $n+1$ affinely independent points, which are the *vertices* of the simplex. A *face* of a geometric simplex is the convex hull of a subset of its vertices. A family C of geometric simplices is a *geometric simplicial complex* if it contains the faces of its members and if whenever two geometric simplices of C intersect, that intersection is a face of both simplices. Let X be a simplicial complex and C a geometric simplicial complex. We say that C is a *geometric realization* of X if there exists a bijection ϕ between the vertices of C and those of X such that a family of vertices in C define a geometric simplex if and only if their images through ϕ define a simplex of X. Just like any graph admits a proper straight-line embedding in $\mathbb{R}^3$, any simplicial complex X with simplices of dimension at most d can be realized by a geometric simplicial complex in $\mathbb{R}^{2d+1}$: map the vertices to any set of points in generic position and define C as the collection of geometric simplices spanned by the images of simplices of X; the generic position of the points ensure that the affine hulls of two subsets of at most d points are disjoint, which in turn implies that C is a geometric simplicial complex.

One can also take a more topological approach, and define a *realization* of a simplicial complex X as a topological space $|X|$ where each d-simplex of X corresponds to a geometric d-simplex (by definition, a geometric (-1)-simplex is empty). We build up the realization of X by increasing dimension. First, create a single point for every vertex (simplex of dimension 0) of X. Then, assuming all the simplices of dimension up to $d-1$ have been realized, consider a d-simplex σ of X. The open lower interval $[0, \sigma)$ is isomorphic to the boundary of the d-simplex by definition; we simply glue a geometric d-simplex to the realization of that boundary. For our purpose, the choice of which types of realizations has little relevance.

Let Δ_n denote some fixed geometric n-simplex in $\mathbb{R}^n$. A *singular n-simplex* of a topological space Γ is a continuous map $\sigma : \Delta_n \to S$. For $n \geqslant 0$, let $C_n(\Gamma)$ be the $\mathbb{Q}$-vector space with basis the set of singular n-simplices of S. An element of $C_n(S)$ is called a *singular n-chain*. The restriction of a singular n-simplex σ to a k-dimensional face f of Δ_n can be seen, via $f \simeq \Delta_k$, as a singular k-simplex of S. We can therefore define face operators, boundary operators, boundaries and cycles similarly as in the simplicial case. This leads to define the *singular homology groups* of the space Γ, and the simplicial homology groups of a simplicial complex X can be shown to be isomorphic to the singular homology group of any realization $|X|$ of X. The analogue of Example B.2 in the singular setting, where X is a d-dimensional ball and $Y \simeq \mathbb{S}^{d-1}$ its boundary, then shows that $H_{d-1}(\mathbb{S}^{d-1})$ is non-trivial.

⋆ ⋆ ⋆

Let X be a simplicial poset and let $V = \{v_1, \ldots, v_\ell\}$ be its vertices, with some fixed order. The simplicial n-chains of X are defined as linear combinations over $\mathbb{Q}$ of simplices of X. If σ is an n-dimensional simplex of X, the lower segment $[0, \sigma]$ is isomorphic to the poset of faces of a standard n-simplex $2^{[n]}$; here we choose the isomorphism so that it preserves the ordering on the vertices. Thus, we get $n+1$-faces $d_i(\sigma) \in X$ (for $i = 0, \ldots, n$), each of dimension $n-1$: namely, $d_i(\sigma)$ is the (unique) face of σ whose vertex set is mapped to $\{0, \ldots, n\} \setminus \{i\}$ by the above isomorphism. Extending the maps d_i by linearity, we get the *face operators* $d_i : C_n(X) \to C_{n-1}(X)$, and let $d = \sum_{i=0}^{n} (-1)^i d_i$ be the boundary operator. The fact that $d \circ d = 0$ follows from the same argument as for simplicial complexes since it is computed inside the vector space generated by $[0, \sigma]$, which is isomorphic to a standard simplex. The definitions of cycles, boundaries and homology groups are then straightforward.

To every simplicial poset X, we can associate a topological space $|X|$, its *realization*, by the same procedure as above: first create a single point for every vertex of X and then, since the open lower interval $[0, \sigma)$ is isomorphic to the boundary of the d-simplex, glue a geometric d-simplex to the realization of that boundary. As for simplicial complexes, the simplicial homology groups of a simplicial poset are isomorphic to the singular homology groups of any of its realizations.

INDEX

PUBLICATIONS CO-AUTHORED BY THE CANDIDATE

[1] B. Aronov, O. Cheong, X. Goaoc, and G. Rote. Lines pinning lines. *Discrete & Computational Geometry*, 45(2):230–260, 2011. (Cited on pages iii, v, 12, 55, 56, 60, 63, 64, 65, 68, and 71.)

[2] D. Attali, O. Devillers, and X. Goaoc. The effect of noise on the number of extreme points. INRIA Research report No. 7134, 2010. (Cited on page iv.)

[3] G. Batog and X. Goaoc. Inflating balls is np-hard. *International Journal of Computational Geometry and Applications*, 2011. To appear.

[4] G. Batog, X. Goaoc, and J. Ponce. Admissible linear map models of linear cameras. In *Proc. 23rd IEEE Conference in Computer Vision and Pattern Recognition (CVPR) 2010*, 2010. (Cited on page iv.)

[5] C. Borcea, X. Goaoc, S. Lazard, and S. Petitjean. On tangents to quadric surfaces. `http://www.arxiv.org/abs/math.AG/0402394`, 2004.

[6] C. Borcea, X. Goaoc, S. Lazard, and S. Petitjean. Common tangents to spheres in $\mathbb{R}^3$. *Discrete & Computational Geometry*, 35(2):287–300, 2006. (Cited on pages 18, 81, and 108.)

[7] C. Borcea, X. Goaoc, and S. Petitjean. Line transversals to disjoint balls. In *Proc. 23rd Ann. Symp. on Computational Geometry*, pages 245–254, 2007. (Cited on pages iii, v, and 22.)

[8] C. Borcea, X. Goaoc, and S. Petitjean. Line transversals to disjoint balls. *Discrete & Computational Geometry*, 1-3:158–173, 2008. (Cited on pages iii, v, 15, 17, 19, 20, 21, and 22.)

[9] H. Brönnimann, O. Devillers, V. Dujmovic, H. Everett, M. Glisse, X. Goaoc, S. Lazard, H.-S. Na, and S. Whitesides. The number of lines tangent to arbitrary polytopes in $\mathbb{R}^3$. In *Proc. 20th ACM Symposium on Computational Geometry (SoCG)*, pages 46–55, 2004.

[10] H. Brönnimann, O. Devillers, V. Dujmovic, H. Everett, M. Glisse, X. Goaoc, S. Lazard, H.-S. Na, and S. Whitesides. Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra. *SIAM Journal of Computing*, 37(2):522–551, 2007.

[11] O. Cheong, X. Goaoc, and A. Holmsen. Hadwiger and Helly-type theorems for disjoint unit spheres in $\mathbb{R}^3$. In *Proc. 21st Ann. Symp. on Computational Geometry*, pages 10–15, 2005. (Cited on pages iii, v, 22, 34, and 46.)

[12] O. Cheong, X. Goaoc, and A. Holmsen. Lower bounds to helly numbers of line transversals to disjoint congruent balls. *Israël Journal of Mathematics*, 2010. To appear. (Cited on pages iii, v, 29, 31, 33, and 34.)

[13] O. Cheong, X. Goaoc, A. Holmsen, and S. Petitjean. Hadwiger and Helly-type theorems for disjoint unit spheres. *Discrete & Computational Geometry*, 1-3:194–212, 2008. (Cited on pages iii, v, 16, 22, 25, 27, 28, 34, and 46.)

[14] O. Cheong, X. Goaoc, and H.-S. Na. Disjoint unit spheres admit at most two line transversals. In *Proc. 11th Annu. European Sympos. Algorithms*, volume 2832 of *Lecture Notes in Computer Science*, pages 127–135, 2003. (Cited on pages iii, v, and 46.)

[15] O. Cheong, X. Goaoc, and H.-S. Na. Geometric permutations of disjoint unit spheres. *Computational Geometry: Theory & Applications*, 30:253–270, 2005. (Cited on pages iii, v, 40, 42, and 46.)

[16] O. Cheong, X. Goaoc, and C. Nicaud. Set systems and families of permutations with small traces. INRIA Research report No. 7154. `http://fr.arxiv.org/abs/0912.2979`, 2009. (Cited on pages iv and 47.)

[17] V. Cortier, X. Goaoc, M. Lee, and H.-S. Na. A note on maximally repeated sub-patterns of a point set. *Discrete Mathematics*, 306(16):1965–1968, 2006.

[18] M. de Berg, X. Goaoc, and A. F. Van der Stappen. A polynomial-time algorithm to design push plans for sensorless parts sorting. In *Proc. Robotics Science and Systems (RSS)*, 2005.

[19] E. Colin de Verdiere, G. Ginot, and X. Goaoc. Helly numbers of acyclic families. `http://arxiv.org/abs/1101.6006`, 2011. (Cited on pages iv, v, 95, 97, 100, 102, and 104.)

[20] J. Demouth, O. Devillers, M. Glisse, and X. Goaoc. Helly-type theorems for approximate covering. In *Proc. 24th ACM Symposium on Computational Geometry (SoCG)*, pages 120–128, 2008. (Cited on page iv.)

[21] J. Demouth, O. Devillers, M. Glisse, and X. Goaoc. Helly-type theorems for approximate covering. *Discrete & Computational Geometry*, 42(3):379–398, 2009. (Cited on page iv.)

[22] J. Demouth and X. Goaoc. Computing direct shadows cast by convex polyhedra. In *European Workshop on Computational Geometry*, 2009. (Cited on page iv.)

[23] O. Devillers, J. Erickson, and X. Goaoc. Empty-ellipse graphs. In *Proc. 19th Symposium on Discrete Algorithms (SODA)*, pages 1249–1257, 2008. (Cited on page iv.)

[24] X. Goaoc. Some discrete properties of the space of line transversals to disjoint balls. In *Non-linear Computational Geometry*, volume 151 of *IMA Volume Series*, pages 51–84. Springer-Verlag, 2009.

[25] X. Goaoc, H.-S. Kim, and S. Lazard. Bounded-curvature shortest path through a sequence of points. INRIA Research report No. 7465. http://hal.inria.fr/inria-00539957/en, 2010. (Cited on page v.)

[26] X. Goaoc, S. Konig, and S. Petitjean. Pinning a line by balls or ovaloids in $\mathbb{R}^3$. *Discrete & Computational Geometry*, 45(2):230–260, 2011. (Cited on pages iii, v, 27, 79, and 80.)

[27] X. Goaoc, J. Kratochvil, Y. Okamoto, C.-S. Shin, A. Spillner, and A. Wolff. Untangling a planar graph. *Discrete & Computational Geometry*, 42(4):542–569, 2009. (Cited on page v.)

[28] X. Goaoc, J. Kratochvil, Y. Okamoto, C.-S. Shin, and A. Wolff. Moving vertices to make drawings plane. In *Proc. 15th International Symposium on Graph Drawing (GD)*, volume 4875 of *Lecture Notes in Computer Science*, pages 101–112, 2007. (Cited on page v.)

OTHER REFERENCES

[AAS99] P. K. Agarwal, B. Aronov, and M. Sharir. Line transversals of balls and smallest enclosing cylinders in three dimensions. *Discrete & Computational Geometry*, 21:373–388, 1999. (Cited on page 109.)

[ABF06] G. Ambrus, A. Bezdek, and F. Fodor. A Helly-type transversal theorem for n-dimensional unit balls. *Archiv der Mathematik*, 86(5):470–480, 2006. (Cited on pages 5 and 23.)

[AK95] N. Alon and G. Kalai. Bounding the piercing number. *Discrete & Computational Geometry*, 13:245–256, 1995. (Cited on page 10.)

[AK05] A. Asinowski and M. Katchalski. Forbidden families of geometric permutations in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 34:1–10, 2005. (Cited on page 46.)

[ALS11] Dominique Attali, André Lieutier, and David Salinas. Efficient Data Structure for Representing and Simplifying Simplicial Complexes in High Dimensions. In *27th Annual Symposium on Computational Geometry 27th Annual Symposium on Computational Geometry*, pages –, 2011. 9 pages. (Cited on page 3.)

[Ame94] N. Amenta. Helly-type theorems and generalized linear programming. *Discrete & Computational Geometry*, 12:241–261, 1994. (Cited on pages 2, 47, and 96.)

[Ame96] N. Amenta. A new proof of an interesting Helly-type theorem. *Discrete & Computational Geometry*, 15:423–427, 1996. (Cited on pages 10, 85, and 89.)

[ARW89] D. Avis, J-M. Robert, and R. Wenger. Lower bounds for line stabbing. *Information processing letters*, 33:59–62, 1989. (Cited on page 47.)

[Asi99] A. Asinowski. Common transversals and geometric permutations. Master thesis, Technion IIT, Haifa, 1999. (Cited on pages 46 and 104.)

[AW10] Gennadiy Averkov and Robert Weismantel. Transversal numbers over subsets of linear spaces. *Adv. Geom.*, 2010. To appear; arXiv.org:1002.0948. (Cited on page 2.)

[Bat08] G. Batog. Événements visuels de polyèdres. Master's thesis, MPRI, ENS Cachan, 2008. (Cited on page iv.)

[Bjö03] A. Björner. Nerves, fibers and homotopy groups. *Journal of Combinatorial Theory, Series A*, 102(1):88–93, 2003. (Cited on pages 3 and 35.)

[BK86] E. Brieskorn and H. Knörrer. *Plane Algebraic Curves*. Birkhaüser Verlag, Basel, 1986. (Cited on page 19.)

[Bor48] K. Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fundamenta Mathematicae*, 35:217–234, 1948. (Cited on pages 3 and 35.)

[Bor06] C. Borcea. Involutive sextics and tangents to spheres. Manuscript, 2006. (Cited on page 19.)

[Bre81] Marilyn Breen. k-dimensional intersections of convex sets and convex kernels. *Discrete Mathematics*, 36(3):233 – 237, 1981. (Cited on page 2.)

[Bre92] Marilyn Breen. Illumination for unions of boxes in $\mathbb{R}^d$. *Proceedings of the American Mathematical Society*, 116:197–202, 1992. (Cited on page 2.)

[Bre97] G. E. Bredon. *Sheaf theory*, volume 170 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1997. (Cited on page 100.)

[Bre03] M. Breen. The dimension of the kernel in an intersection of starshaped sets. *Archiv der Mathematik*, 81:485–490, 2003. (Cited on page 2.)

[BT82] R. Bott and L. W. Tu. *Differential forms in algebraic topology*, volume 82 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1982. (Cited on page 3.)

[BTRBI79] A. Ben-Tal, E. Rosinger, and A. Ben-Israel. A helly type theorem and semiinfinite programming. In *Constructive Approaches to Mathematical Models*, pages 127–135. Academic Press, 1979. (Cited on page 3.)

[CHY10] Otfried Cheong, Jae-Soon Ha, and Jungwoo Yang. Personal communication, 2010. (Cited on pages 39 and 46.)

[Cla95] Kenneth L. Clarkson. Las vegas algorithms for linear and integer programming when the dimension is small. *J. ACM*, 42:488–499, 1995. (Cited on page 96.)

[Dan57] L. Danzer. Über ein Problem aus der kombinatorischen Geometrie. *Archiv der Mathematik*, 1957. (Cited on page 4.)

[Dem08] J. Demouth. Événements visuels de convexes et limites d'ombres. PhD thesis, Université Nancy 2, 2008. (Cited on page iv.)

[DGK63] L. Danzer, B. Grünbaum, and V. Klee. Helly's theorem and its relatives. In V. Klee, editor, *Convexity*, Proc. of Symposia in Pure Math., pages 101–180. Amer. Math. Soc., 1963. (Cited on pages 2 and 6.)

[Dur00] F. Durand. A multidisciplinary survey of visibility. In *ACM Siggraph Course Notes: Visibility, Problems, Techniques, and Applications*, 2000. (Cited on page 113.)

[Eck93] J. Eckhoff. Helly, Radon and Caratheodory type theorems. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Convex Geometry*, pages 389–448. North Holland, 1993. (Cited on pages 2 and 6.)

[Edm01] Jack Edmonds. Redundancy and helly. *European Journal of Combinatorics*, 22(5):679 – 685, 2001. (Cited on page 2.)

[EN09] J. Eckhoff and K.-P. Nischke. Morris's pigeonhole principle and the Helly theorem for unions of convex sets. *Bulletin of the London Mathematical Society*, 41:577–588, 2009. (Cited on pages 10, 85, and 92.)

[EW89] P. Egyed and R. Wenger. Stabbing pairwise disjoint translates in linear time. In *Proc. 5th Symposium on Computational Geometry*, pages 364–369, 1989. (Cited on page 47.)

[Far09] Benson Farb. Group actions and helly's theorem. *Advances in Mathematics*, 222(5):1574 – 1588, 2009. (Cited on page 3.)

[Flø11] Gunnar Fløystad. The colorful helly theorem and colorful resolutions of ideals. *Journal of pure and applied Algebra*, 2011. To appear. (Cited on page 3.)

[Gar92] B. Gartner. A subexponential algorithm for abstract optimization problems. In *Proc. 33rd Annual Symposium on Foundations of Computer Science (FoCS)*, pages 464–472, 1992. (Cited on page 47.)

[GJ99] P. G. Goerss and J. F. Jardine. *Simplicial homotopy theory*, volume 174 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1999. (Cited on pages 98 and 99.)

[GM61] B. Grünbaum and T.S. Motzkin. On components in some families of sets. *Proceedings of the American Mathematical Society*, 12:607–613, 1961. (Cited on pages 10, 85, 86, and 95.)

[GM93] V. Goryunov and D. Mond. Vanishing cohomology of singularities of mappings. *Compositio Math.*, 89(1):45–80, 1993. (Cited on page 89.)

[GM10] Stéphane Gaubert and Frédéric Meunier. Carathéodory, helly and the others in the max-plus world. *Discrete & Computational Geometry*, 43:648–662, 2010. (Cited on page 2.)

[God73] R. Godement. *Topologie algébrique et théorie des faisceaux*. Hermann, Paris, 1973. Troisième édition revue et corrigée, Publications de l'Institut de Mathématique de l'Université de Strasbourg, XIII, Actualités Scientifiques et Industrielles, No. 1252. (Cited on page 100.)

[GP02] J. E. Goodman and R. Pollack. Geometric transversal theory. In *Encyclopaedia of Mathematics*. Springer-Verlag, Heidelberg, Germany, 2002. `http://eom.springer.de/g/g130050.htm`. (Cited on page 6.)

[GPW93] J. E. Goodman, R. Pollack, and R. Wenger. Geometric transversal theory. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 163–198. Springer-Verlag, Heidelberg, Germany, 1993. (Cited on page 6.)

[Gre67] M. J. Greenberg. *Lectures on Algebraic Topology*. Benjamin, Reading, MA, 1967. (Cited on pages 89, 100, and 101.)

[Grü58] B. Grünbaum. On common transversals. *Archiv der Mathematik*, 9:465–469, 1958. (Cited on page 4.)

[GS05] E. Greenstein and M. Sharir. The space of line transversals to pairwise disjoint balls in $\mathbb{R}^3$. manuscript, 2005. (Cited on page 5.)

[GS08] S. Gaubert and S. Sergeev. Cyclic projectors and separation theorems in idempotent convex geometry. *Journal of Mathematical Sciences*, 155:815–829, 2008. (Cited on page 2.)

[Had] H. Hadwiger. Problem 107. *Nieuw Archiv Wiskunde*, (3)4:57, 1956; Solution. *Wiskundige Opgaven*, 20:27–29, 1957. (Cited on pages 5 and 23.)

[Had57] H. Hadwiger. Über eibereiche mit gemeinsamer treffgeraden. *Portugal Math.*, 6:23–29, 1957. (Cited on pages 27 and 34.)

[Had63] A. Glen Haddock. Some theorems related to a theorem of e. helly. *Proceedings of the American Mathematical Society*, 14:636–637, 1963. (Cited on page 3.)

[Hal04] Nir Halman. On the power of discrete and of lexicographic helly-type theorems. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 463–472. IEEE Computer Society, 2004. (Cited on page 3.)

[Hat02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. (Cited on page 115.)

[Hel23] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresbericht Deutsch. Math. Verein.*, 32:175–176, 1923. (Cited on pages 1 and 12.)

[Hel30] E. Helly. Über Systeme von abgeschlossenen Mengen mit gemeinschaftlichen Punkten. *Monaths. Math. und Physik*, 37:281–302, 1930. (Cited on page 3.)

[HKL03] A. Holmsen, M. Katchalski, and T. Lewis. A Helly-type theorem for line transversals to disjoint unit balls. *Discrete & Computational Geometry*, 29:595–602, 2003. (Cited on pages 5 and 23.)

[HM04] A. Holmsen and J. Matoušek. No Helly theorem for stabbing translates by lines in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 31:405–410, 2004. (Cited on page 61.)

[HP94] W. V. D. Hodge and D. Pedoe. *Methods of Algebraic Geometry*, volume 2, chapter XIV: Grassmann varieties. Cambridge University Press, 2 edition, 1994. (Cited on page 114.)

[HXC01] Y. Huang, J. Xu, and D. Z. Chen. Geometric permutations of high dimensional spheres. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 244–245, 2001. (Cited on page 46.)

[HXC04] Y. Huang, J. Xu, and D. Z. Chen. Geometric permutations of high dimensional spheres. *Computational Geometry: Theory & Applications*, 29:47–60, 2004. (Cited on page 46.)

[Jan10] J. Jang. Computing direct shadows cast among disjoint convex polyhedra. Master's thesis, KAIST University, Korea, 2010. (Cited on page iv.)

[Kat86] M. Katchalski. A conjecture of Grünbaum on common transversals. *Math. Scand.*, 59(2):192–198, 1986. (Cited on page 104.)

[KLL87] M. Katchalski, T. Lewis, and A. Liu. Geometric permutations of disjoint translates of convex sets. *Discrete Mathematics*, 65:249–259, 1987. (Cited on page 104.)

[KM06] G. Kalai and R. Meshulam. Intersections of leray complexes and regularity of monomial ideals. *Journal of Combinatorial Theory Series A*, 113:1586–1592, 2006. (Cited on pages 3 and 100.)

[KM08] G. Kalai and R. Meshulam. Leray numbers of projections and a topological helly-type theorem. *Journal of Topology*, 1:551–556, 2008. (Cited on pages 3, 10, 11, 85, 88, 95, and 100.)

[KMY92] David Kirkpatrick, Bhubaneswar Mishra, and Chee-Keng Yap. Quantitative steinitz's theorems with applications to multifingered grasping. *Discrete & Computational Geometry*, 7:295–318, 1992. (Cited on page 34.)

[Kol91] Krzysztof Kolodziejczyk. Generalised helly and radon numbers. *Bulletin of the Australian Mathematical Society*, 43:429–437, 1991. (Cited on page 2.)

[KR06] Umed H. Karimov and Dusan Repovs. On the topological helly theorem. *Topology and its Applications*, 153(10):1614 – 1621, 2006. (Cited on page 35.)

[KS90] M. Kashiwara and P. Schapira. *Sheaves on manifolds*, volume 292 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1990. (Cited on page 100.)

[KSZ03] M. Katchalski, S. Suri, and Y. Zhou. A constant bound for geometric permutations of disjoint unit balls. *Discrete & Computational Geometry*, 29:161–173, 2003. (Cited on page 46.)

[KV01] Matthew J. Katz and Kasturi R. Varadarajan. A tight bound on the number of geometric permutations of convex fat objects in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 26:543–548, 2001. (Cited on page 35.)

[Lay72] Steven R. Lay. Separation by cylindrical surfaces. *Proceedings of the American Mathematical Society*, 36:224–228, 1972. (Cited on page 2.)

[LS09] Michael Langberg and Leonard J. Schulman. Contraction and expansion of convex sets. *Discrete & Computational Geometry*, 42:594–614, 2009. (Cited on page 3.)

[Mas01] M. T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, Cambridge, USA, 2001. (Cited on page 34.)

[Mat97] J. Matoušek. A Helly-type theorem for unions of convex sets. *Discrete & Computational Geometry*, 18:1–12, 1997. (Cited on pages 3, 10, and 104.)

[May92] J. Peter May. *Simplicial objects in algebraic topology*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992. Reprint of the 1967 original. (Cited on pages 98 and 99.)

[MB62] J. Milnor and M. Barratt. An example of anomalous singular homology. *Proceedings of the American Mathematical Society*, 13:293–297, 1962. (Cited on page 95.)

[Meg90] N. Megiddo. On the complexity of some geometric problems in unbounded dimension. *Journal of Symbolic Computation*, 10:327–334, 1990. (Cited on page 47.)

[Meg01] G. Megyesi. Lines tangent to four unit spheres with coplanar centers. *Discrete & Computational Geometry*, 26:493–497, 2001. (Cited on pages 81 and 108.)

[MNP90] X. Markenscoff, L. Ni, and C. H. Papadimitriou. The geometry of grasping. *Internat. J. Robot. Res.*, 9:61–74, 1990. (Cited on page 34.)

[Mor73] H.C. Morris. *Two pigeon Hole Principles and Unions of Convexly Disjoint Sets*. Phd thesis, California Institute of Technology, 1973. (Cited on pages 10 and 85.)

[MPT01] I.G. Macdonald, J. Pach, and T. Theobald. Common tangents to four unit balls in $\mathbb{R}^3$. *Discrete & Computational Geometry*, 26:1–17, 2001. (Cited on pages 81 and 108.)

[MS05] G. Megyesi and F. Sottile. The envelope of lines meeting a fixed line and tangent to two spheres. *Discrete & Computational Geometry*, 33(4):617–644, 2005. arXiv math.AG/0304346. (Cited on pages 81 and 108.)

[MSS87] B. Mishra, J. T. Schwartz, and M. Sharir. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2:541–558, 1987. (Cited on page 34.)

[MST03] G. Megyesi, F. Sottile, and T. Theobald. Common transversals and tangents to two lines and two quadrics in $\mathbb{P}^3$. *Discrete & Computational Geometry*, 30:543–571, 2003. (Cited on pages 81 and 108.)

[MSW96] J. Matousek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. *Algorithmica*, 16:498–516, 1996. (Cited on page 96.)

[MT04] A. Marcus and G. Tardos. Excluded permutation matrices and the Stanley–Wilf conjecture. *Journal of Combinatorial Theory. Series A*, 107(1):153–160, 2004. (Cited on page 47.)

[MT08] J. Matoušek and M. Tancer. Dimension gaps between representability and collapsibility. *Discrete & Computational Geometry*, online first, 2008. (Cited on page 3.)

[Pet71] C. M. Petty. Equilateral sets in minkowski spaces. *Proceedings of the American Mathematical Society*, 29:369–374, 1971. (Cited on page 2.)

[PW01] H. Pottmann and J. Wallner. *Computational Line Geometry*. Springer, 2001. (Cited on pages 75, 81, and 114.)

[RKS10] N. Rubin, H. Kaplan, and M. Sharir. Improved bounds for geometric permutations. In *Proc. 51st IEEE Symp. on Foundations of Computer Science*, pages 355–364, 2010. (Cited on page 35.)

[Rob42] C. V. Robinson. Spherical theorems of Helly type and congruence indices of spherical caps. *American Journal of Mathematics*, 64:260–272, 1942. (Cited on page 54.)

[Sei91] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete & Computational Geometry*, 6:423–434, 1991. (Cited on page 96.)

[SMS00] S. Smorodinsky, J. S. B. Mitchell, and M. Sharir. Sharp bounds on geometric permutations for pairwise disjoint balls in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 23:247–259, 2000. (Cited on page 35.)

[Spa66] E. H. Spanier. *Algebraic topology*. McGraw-Hill Book Co., New York, 1966. (Cited on page 100.)

[ST08] F. Sottile and T. Theobald. Line problems in non-linear computational geometry. In J. E. Goodman, J. Pach, and R. Pollack, editors, *Computational Geometry - Twenty Years Later*, pages 411–432. AMS, 2008. (Cited on pages 81 and 108.)

[SW92] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. In *Proc. 9th Sympos. on Theo. Aspects of Comp. Science*, pages 569–579, 1992. (Cited on pages 2 and 96.)

[Swa03] Konrad J. Swanepoel. Helly-type theorems for homothets of planar convex curves. *Proceedings of the American Mathematical Society*, 131:921–932, 2003. (Cited on page 2.)

[Tan10] Martin Tancer. A counterexample to wegner's conjecture on good covers. arxiv.org/abs/1008.1895, 2010. (Cited on page 95.)

[Theo02] T. Theobld. An enumerative geometry framework for algorithmic line problems in $\mathbb{R}^3$. *SIAM Journal of Computing*, 31:1212–1228, 2002. (Cited on pages 81 and 108.)

[Topo06] V. Toponogov. *Differential Geometry of Curves and Surfaces: A Concise Guide*. Birkhäuser, 2006. (Cited on page 21.)

[Tve89] H. Tverberg. Proof of Grünbaum's conjecture on common transversals for translates. *Discrete & Computational Geometry*, 4:191–203, 1989. (Cited on pages 4 and 104.)

[Vin35] P. Vincensini. Figures convexes et variétés linéaires de l'espace euclidien à n dimensions. *Bull. Sci. Math.*, 59:163–174, 1935. (Cited on page 22.)

[Weg75] G. Wegner. d-collapsing and nerves of families of convex sets. *Archiv der Mathematik*, 26:317–321, 1975. (Cited on page 95.)

[Wen90] R. Wenger. Upper bounds on geometric permutations for convex sets. *Discrete & Computational Geometry*, 5:27–33, 1990. (Cited on page 35.)

[Wen04] R. Wenger. Helly-type theorems and geometric transversals. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete & Computational Geometry*, chapter 4, pages 73–96. CRC Press LLC, Boca Raton, FL, 2nd edition, 2004. (Cited on pages 2 and 6.)

[ZS01] Y. Zhou and S. Suri. Shape sensitive geometric permutations. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 234–243, 2001. (Cited on page 46.)