Tropical bisectors and diameters of simplicial complexes

vorgelegt von
M. Sc.
Francisco Criado Gallart

von der Fakultät II – Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
Dr.rer.nat.

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender:	Prof. Dr. Jochen Blath
Gutachter:	Prof. Dr. Michael Joswig
Gutachter:	Prof. Dr. Günter Rote
Gutachter:	Prof. Dr. Francisco Santos Leal

Tag der wissenschaftlichen Aussprache: 30.11.2020

Berlin 2021

Summary

Tropical geometry is a discrete analogue of algebraic geometry where addition is replaced by taking the minimum (or maximum) and multiplication is replaced by addition. It is an active research field where tropical analogues of geometric concepts, like linear spaces, varieties, polytopes or volume, are in different stages of development. The interest of these analogues is twofold: their combinatorial nature helps us understand difficult geometric objects and develop algorithms for them, and at the same time they bring geometric language to the field of mathematical optimization.

The application of tropical geometry to linear programming has been introduced in [5], where tropical ideas were used to prove that log-barrier interior point linear programming algorithms, known to run in weakly polynomial time, cannot run in strongly polynomial time. In a similar spirit, we are interested in the study of the tropical distance and in particular how it relates to existing tropical linear programming algorithms.

The problem of algorithmically solving tropical linear programs is particularly interesting because it lies in the classes of NP and co-NP problems, but no polynomial time algorithm for it is known yet. The problem has an interpretation in the context of metric tropical geometry, but it is still unknown how to use this geometric information for tropical linear feasibility algorihtms.

For this reason we introduce here the study of tropical bisectors and Voronoi diagrams, as well as algorithms for their computation. Even though our study focuses on these objects on their own, they will be necessary for understanding tropical analogues for the ellipsoid method or the Kaczmarz iteration algorithm. It is relevant to develop such analogues, because both of these algorithms run in (weakly) polynomial time for classical linear programming.

The rest of the thesis looks at the diameter of simplicial complexes, a combinatorial generalization of polytopes. Both classical and tropical simplex methods for linear programming involve exploring the dual graph of some simplicial sphere. Thus, studying the diameter of such complexes is relevant for the understanding of the simplex method's computational complexity. We present new constructions for abstract simplicial complexes with asymptotically large diameter, close to the theoretical upper bound, using a probabilistic construction. We also construct small simplicial spheres with large diameter using a computational heuristic approach, and we develop the topological analogue of the prismatoids used by Santos to disprove the Hirsch conjecture.

Zusammenfassung

Tropische Geometrie ist ein diskretes Analogon zu algebraischer Geometrie, bei dem Addition durch Minimumsbildung (oder Maximumsbildung) und Multiplikation durch Addition ersetzt wird. Sie ist ein aktives Forschungsfeld; tropische Analoga zu geometrischen Konzepten wie Volumina, Polytopen, linearen Räumen oder Varietäten sind in unterschiedlichen Stadien der Entwicklung. Diese Analoga sind in zweierlei Hinsicht interessant: Zum einen hilft uns ihre kombinatorische Natur, komplizierte geometrische Objekte zu verstehen und Algorithmen für sie zu entwickeln. Zum anderen ermöglichen sie die Verwendung geometrischer Sprache für mathematische Optimierung.

Anwendungen tropischer Geometrie auf lineare Programmierung wurden erstmals in [5] untersucht. Dort werden tropische Ideen verwendet, um zu beweisen, dass Innere-Punkte-Methoden mit logarithmischer Barrier-Funktion, die bekanntermaßen in schwach polynomineller Zeit terminieren, nicht in stark polynomineller Zeit terminieren können. Diesem Ansatz folgend untersuchen wir hier die tropische Abstandsfunktion auch im Zusammenhang mit bestehenden Algorithmen zur tropischen linearen Programmierung.

Das Problem der algorithmischen Lösung tropischer linearer Programme ist besonders interessant, da es im Schnitt der Klassen NP und Co-NP liegt, aber noch kein Lösungsalgorithmus mit polynomineller Laufzeit bekannt ist. Es lässt sich im Kontext metrischer tropischer Geometrie interpretieren; bisher ist allerdings unklar, wie sich diese geometrische Information in Algorithmen zur Bestimmung der Lösbarkeit tropischer linearer Programme einbinden lässt.

Aus diesem Grund stellen wir hier eine Untersuchung von tropischen Voronoi-Diagrammen, tropischen Bisektoren und Algorithmen zu ihrer Berechnung vor. Unsere Arbeit legt den Fokus auf die Untersuchung dieser Objekte an sich; sie werden aber auch für das Verständnis tropischer Analoga zur Ellipsoidmethode oder zur Kaczmarz-Methode notwendig sein. Die Entwicklung solcher Analoga ist relevant, da beide Algorithmen für klassische lineare Programme (schwach) polynominelle Laufzeit aufweisen.

Der Rest der Arbeit befasst sich mit dem Durchmesser von Simplizialkomplexen, einer kombinatorischen Verallgemeinerung von Polytopen. Sowohl klassische als auch tropische Simplex-Methoden für lineare Programmierung beinhalten die Untersuchung des dualen Graphen einer simplizialen Sphäre. Die Untersuchung des Durchmessers solcher Komplexe ist daher für das Verständnis der Komplexität dieser Methoden relevant. Wir präsentieren neue Konstruktionen abstrakter Simplizialkomplexe mit asymptotisch großem Durchmesser nahe an der theoretischen oberen Schranke, unter Verwendung einer probabilistischen Konstruktion. Wir konstruieren weiterhin, mithilfe eines heuristischen Ansatzes, kleine simpliziale Sphären mit großem Durchmesser, und entwickeln das topologische Analogon zu den Prismatoiden, die Santos verwendet um die Hirsch-Vermutung zu widerlegen.

Acknowledgements

First I would like to thank my advisors Michael Joswig and Francisco Santos for their scientific guideance. Without their persistent help, this project would have never been realized.

I wish to acknowledge the funding and help from the Berlin Mathematical School. I always felt like I could count with them for so many things.

I would like also to thank my research group for contributing to such a great work environment. In particular to my coauthor Andrew Newman and to Antje, whose organizational help was always crucial, as well as Benjamin and Lars for their technical help.

Thanks also to Aitor, David, Jaime, Manuel, Robert, Olivia and Sofia for their proofreading, and to Georg and Tarek for helping me think about tropical linear programming in a different light.

And finally thanks to my friends and family for years of support, without which this thesis would have not be finished. Ander, Juanín, René, Stephen, Pedro, Irene, mom, dad, and many others. Thank you so much.

Contents

Introduction

A *linear program* is an optimization problem where the objective is to minimize (or maximize) a linear objective function while satisfying a set of linear constraints over $\mathbb{R}^d$. This is, it is the problem of optimizing a linear function over a polyhedron. Linear programming is one of the most elementary problems in the field of mathematical optimization, and yet, it is unknown whether strongly polynomial time algorithms exist to solve these problems; this question is actually Problem 9 in Smale's list of "Mathematical problems for the 21st century" [65].

Roughly speaking, an algorithm runs in *strongly polynomial time* if it runs in polinomial time in both the *arithmetic model of computation* and in the *bit model of computation*. In the former, each arithmetic operation takes time $O(1)$, whereas in the latter, only bit operations take constant time and arithmetic operations must be implemented with suitable algorithms. In contrast, an algorithm runs in *weakly polynomial time* if it runs in polynomial time in the bit model of computation. This is commonly known as a *polynomial time algorithm*. Observe that, in practice, the arithmetic model is a better representation of time complexity if approximate solutions suffice and we use fixed-precision numbers.

One of the earliest references to the study of linear optimization is due to Fourier, who in 1827 published a method for solving systems of linear inequalities [26]. This problem of determining the *linear feasibility* of a set of linear inequalities is equivalent to approximately solving linear programs: if we know how to solve one problem in weakly polynomial time, we can solve the other one. One direction is trivial, and the other can be shown with binary search.

The first general purpose linear programming solver was the *simplex method*, developed in 1947 by George Dantzig, originally to solve planning problems for the U.S. Air Force [18]. The simplex method is still one of the most widely used classes of algorithms for linear programming because of its practical efficiency. However, surprisingly little is known about its theoretical efficiency. In fact, we do not know if it is weakly polynomial or not.

The simplex method works by exploring the *edge graph* of a polyhedron. It starts at some vertex, then moves from vertex to vertex of the polytope until the current vertex being examined can be proven to be a global optimum.

Many different strategies have been suggested to define exactly how to choose the next vertex to explore. These are called the *pivot rules*. Currently we know of no pivot rule achieving a number of steps polynomial on the number of facets and dimension. Furthermore, if the diameter of the edge graph of a polytope could grow superpolynomially on its number of facets and dimension, then no such pivot rule would exist at all.

This motivates the study of the diameter of the edge graph of simple polyhedra. It begins with a letter from Warren M. Hirsch to Dantzig in 1957, where he conjectures that the diameter of a polyhedron with n facets in dimension d is at most $n-d$. Nowadays we know that this is not the case; the unbounded case was solved in 1967 by Klee and Walkup [46], and the bounded case by Santos in 2012 [63]. However, Santos' construction still produces polyhedra with diameter linear in $(n-d)$. So even with the Hirsch conjecture refuted, we still do not know if this diameter can grow superpolynomially in n and d, which would be relevant for the complexity of the simplex method.

The first (weakly) polynomial time algorithm for linear programming was the *ellipsoid method*, an older algorithm for general purpose convex optimization. Leonid Khachiyan studied its application to the class of linear programming problems in 1979, and proved that it runs in weakly polynomial time [44]. The simplex method used the combinatorial structure of the feasible region, but the ellipsoid method uses only geometric information.

For the linear feasibility problem, it works by covering the feasible region with an ellipsoid that gradually reduces its volume as the algorithm runs. At each iteration we evaluate the inequalities at the ellipsoid's centroid. Either we find that this point is a solution to all of them, or there is a violated inequality. Using this linear inequality, we compute a new ellipsoid whose volume is smaller than the original, and is still guaranteed to cover the feasible region. Furthermore, we can show that the ratio of these volumes is big enough to prove fast convergence. Additional lemmas give us initial ellipsoids that are guaranteed to cover the feasible region, as well as a threshold for the volume of the ellipsoid at which we stop.

The ellipsoid method is mostly of theoretical interest since it requires square root computations, forcing the algorithm to double the precision of intermediate values. But it inspired many other algorithms, like the Yamnitsky–Levin algorithm, which uses a simplex instead of an ellipsiod, thus reducing the precision required to run [69]. In general, these algorithms only require a *separation oracle*: a function that tells us if a query point is in the polyhedron, or gives us a linear space separating the query from the solution set.

The next breakthrough came in 1984, with the first interior point method running in weakly polynomial time, by Narenda Karmakar [43]. These algorithms iteratively evaluate a combination of both the objective function and a barrier function, which prevents the point from getting out of bounds from the polytope. Karmakar's algorithm is the first weakly polynomial time algorithm to be usable in practice, with a practical running time similar to that of the simplex method in some cases.

But these interior point methods are not strongly polynomial time algorithms, as proved by Allamigeon et al. in 2018 [5]. In this article, *tropical geometry* was used to construct a linear program such that the trajectory of the interior point method is twisted, requiring a number of iterations superpolinomial in n and d (but still polynomial in the number of bits of the input). This article was an initial motivation for this thesis, which contains elements of both linear programming and tropical geometry.

Tropical geometry is, at the same time, a discrete analogue of algebraic geometry and an appropriate language for describing discrete optimization problems for graphs and scheduling. Thus, it connects two apparently separate mathematical disciplines. A proper introduction to tropical geometry will be presented in Section 1.3, but for now it suffices to say that its tools enable us to understand complex geometric objects. This, in turn, allows for better algorithms for known geometric problems, and also for the development of extremal examples. At the same time, it brings a geometric language to a discrete optimization context. And, as is usually the case in mathematics, tropical geometry eventually developed its own interesting questions and problems.

We are particularly interested in *tropical linear feasibility* (and *tropical linear programs*), where the constraints (and objective function for tropical linear programs) are tropical linear functions. They arise in network games, and are particularly interesting for and/or scheduling problems [57]. Most importantly, tropical linear programs are in the class of NP and co-NP problems [2], but no polynomial time algorithm is known for them.

How could we design such an algorithm? This thesis explores problems related to two interesting possibilities.

The ellipsoid method was the first polynomial time algorithm for classical linear feasibility, so it is a good candidate for a tropical version. It requires three things: a bounding region in which to contain the potential solution, a notion of volume for proving convergence, and a separation oracle to gain information from each failed attempt.

The bounding region in this case is most naturally a particular type of tropical polyhedron, a *polytrope*. Polytropes are special in that they are simultaneously tropical polytopes and classical polytopes. Classically, they have $d(d+1)$ facets with normal vectors of the form $e_i - e_j \quad \forall i, j \in \{1, \dots, d+1\}$. They are simultaneously the most natural tropical version of a classical simplex and a classical ellipsoid, which make them very good candidates for a tropical ellipsoid (or Yamnitsky–Levin) variant.

A tropical analogue for volume was recently introduced in [37] and [51]. The idea here is to adapt *Ehrhart theory* to the tropical context. Classically, one could define the volume of a set as the limit, as $\varepsilon \to \infty$, of the number of points inside the set in an integer ε-lattice, multiplied by ε^d. Tropical volume uses the same definition, but a *tropical integer* will be not the same thing as a classical integer; a tropical integer is a number of the form $\log(n)$ where $n \in \mathbb{Z}_{\geq 0}$. This definition could be the correct way to describe the bit complexity of a tropical number for the purposes of complexity analysis.

What remains is the separation oracle. The *Shapley operator*, an important function in the context of tropical linear feasibility, could provide one. The Shapley operator takes a point in the ambient space and maps it to another point that, intuitively, is "closer" to being a solution.

Current algorithms for tropical linear feasibility work by iterating the Shapley operator over a point, and averaging the trajectory. This trajectory eventually orbits around a fixed point, which is guaranteed to exist. It is known that a fixed point of this operator gives either a positive or a negative certificate for the tropical linear feasibility problem. Then, the average of the trajectory arbitrarily approaches a fixed point, which gives the certificate.

What these algorithms are missing is the *metric* information the Shapley operator gives. Let us consider the *tropical distance*, defined by

$$\operatorname{dist}(a, b) = \max_{i \in [d+1]} (a_i - b_i) - \min_{j \in [d+1]} (a_j - b_j).$$

Then, the image of a point under the Shapley operator is closer to the set of fixed points in this metric than the original point.

This suggests a separation oracle, except instead of giving us a linear space as a separating object, it gives us the *tropical bisector* between the point and its image. In classical geometry bisectors are also linear spaces, but this does not happen anymore in the tropical setting. If we want to adapt the ellipsoid method to tropical linear feasibility, we need to understand these tropical bisectors and we need to look at the tropical metric from a computational perspective.

This is exactly why we introduce the study of tropical bisectors in Chapter 2. Bisectors for polyhedral norms are already widely studied, but the combinatorial properties of the tropical ball gives us a lot of information and structure that we can use. We study their combinatorics, as well as the topology of bisectors of an arbitrary number of points. We classify and characterize their types, and we study their complexity. Finally, we study the computational problem of constructing bisectors using tools inspired by classical computational geometry.

Another possibility would be to study the complexity of the tropical analogue of the simplex method [4]. The *tropical simplex method* explores the edge graph of a tropical polyhedron instead of a classical one. In both cases, we may look at just the graphs of simple polyhedra, which have exactly d edges coinciding in each vertex. We can consider just these because every non-simple polyhedron can be perturbed into a simple one, which will approximate the original. The edge graphs of bounded simple polyhedra are equivalent (by duality) to the *dual graphs* of *simplicial polytopes*, that is, bounded polyhedra where each facet is a simplex.

In the next two chapters we look at the dual diameter of simplicial complexes, the combinatorial generalization of simplicial polytopes, and simplicial spheres, simplicial complexes homeomorphic to $\mathbb{S}_{d-1}$. The maximum dual diameter of simplicial spheres is also unknown, but every upper bound known for diameters of simplicial polytopes applies to simplicial spheres and viceversa, so we believe these diameters may be asymptotically close.

Chapter 3 constructs abstract simplicial complexes with large diameter using a probabilistic construction. Here we do not restrict the topological nature of the complex, so the diameter grows superpolynomially in n and d. The chapter improves a construction of the author (See [16]) at the expense of no longer being deterministic. While the previous construction's diameter was an exponential factor of $\Theta(e^d)$ away from the trivial upper bound, our new construction is only a factor of $\Theta(d^2)$ away from this upper bound. We also apply the result to pseudo-manifolds, a class of complexes between pure simplicial complexes and combinatorial manifolds. For these, the new gap between upper and lower bounds is $\Theta(d^3)$.

Finally, Chapter 4 constructs small simplicial spheres whose diameter exceed the $n-d$ bound given by the Hirsch Conjecture. We adapt Santos' construction of a polytopal counterexample to the topological setting, and we use a simulated annealing metaheuristic approach to reduce the size of these complexes as much as possible while preserving their topology and non-Hirsch-ness. The main result are four combinatorially different simplicial spheres of dimension 8 and with 18 vertices that violate the Hirsch bound. If they turn out to be also simplicial polytopes (which is an open question at this moment) they would produce 9-dimensional non-Hirsch polytopes with 18 facets, while the smallest non-Hirsch polytopes known to date are 20-dimensional and they have 40 facets.

Summing up, in this thesis we address various problems at the interplay of optimization, computational geometry and topological combinatorics. Our motivations come from understanding the complexity of linear programming, particularly in the tropical version. We believe that tropical linear programming has potential for becoming an important tool if efficient algorithms are found for it, and that our results may contribute to their development.

Chapter 1

Preliminaries

1.1 Polytopes, polyhedra and simplicial complexes

1.1.1 Convex hulls and cones

We begin with the basic definitions for the field of polytope theory. Let $X = \{x_1, \dots, x_n\} \subseteq \mathbb{R}^d$ be a finite set of point (or vectors) in the d-dimensional Euclidean affine space (or vector space). Then,

- The *vector span* of X is $\text{span}(x_1, \dots, x_n) = \{\sum_{i=0}^n \lambda_i x_i \mid \lambda \in \mathbb{R}^n\}$.
- The *affine span* of X is $\text{affine}(x_1, \dots, x_n) = \{\sum_{i=0}^n \lambda_i x_i \mid \lambda \in \mathbb{R}^n, \sum_{i=0}^n \lambda_i = 1\}$.
- The *cone* of X is $\text{cone}(x_1, \dots, x_n) = \{\sum_{i=0}^n \lambda_i x_i \mid \lambda \in \mathbb{R}^n, \lambda \geq 0\}$.
- The *convex hull* of X is $\text{conv}(x_1, \dots, x_n) = \{\sum_{i=0}^n \lambda_i x_i \mid \lambda \in \mathbb{R}^n, \lambda \geq 0, \sum_{i=0}^n \lambda_i = 1\}$.

The first two definitions are standard theory in basic linear algebra. The other two are the nonnegative versions of them.

If X is a finite set, then the cone of X is called a *polyhedral cone*, and the convex hull of X is called a *polytope*.

Intuitively, polytopes are bounded convex subsets of the affine space with a piecewise linear boundary. This is why the notion of *polyhedron* is very relevant too. A polyhedron (plural: polyhedra) is a subset of $\mathbb{R}^d$ defined by linear inequalities, as $Ax \geq b$. A polyhedron defined by just one linear inequality is called a *linear half space*. The *dimension* of a polytope or polyhedron is the dimension of the lowest-dimensional affine subspace that contains it.

The first foundational result is the Minkowski-Weyl theorem:

Theorem 1.1.1 (Minkowski-Weyl theorem [70]). *The classes of bounded polyhedra and of polytopes are exactly the same.*

This *double description* is key for understanding the geometric and combinatorial properties of polytopes.

1.1.2 Faces of a polyhedron

We proceed now to study the combinatorial structure of polytopes. The study for polyhedra is similar, but it adds an extra layer of complexity that is not required for the purposes of this work.

Let P be a polytope, and $v \in R^d$ be a nonzero vector. Consider the subset of P that maximizes the linear functional $f_v(x) = v^t x$, if it exists. This subset is necessarily a lower dimensional polytope, as it is the intersection of P with some linear hyperplane. Furthermore, every point in the boundary of P maximizes some linear functional in this way, because every boundary point is in the border of at least one of the defining inequalities. This means that we can cover the boundary of P with such polytopes. These polytopes that maximize a linear function are called the *faces* of P. The set of faces of a polytope form a partially ordered set by inclusion.

Faces of a polytope are classified by their dimension. Some of these k-dimensional faces have special names:

Facets for $k = d - 1$.

Ridges for $k = d - 2$.

Edges for $k = 2$.

Vertices for $k = 1$.

The number of faces of each type is interesting for combinatorial purposes. This is why they are compiled into the *f-vector*. The f-vector of a polytope is a $(d+1)$ vector with natural entries where f_i is the number of faces of dimension i, for $i = 0 \ldots, d-1$.

Observe that the combinatorial complexity of a polytope with few vertices can be exponential. For example, the *crosspolytope*, is the generalization of the octahedron, defined as $\mathrm{conv}(\{\pm e_i | i \in [d]\})$. This polytope has $2d$ vertices but 2^d facets, one for each assignment of signs to the canonical base.

The upper bound theorem gives the maximum number of faces of each type a polytope can have:

Theorem 1.1.2 (Upper bound theorem)**.** *For any d-polytope P with n vertices, and every $i = 0, \ldots, d-1$*

$$f_i(P) \leq f_i(\Delta(n, d)).$$

Where $\Delta(n, d)$ is the cyclic polytope *defined as the convex hull of n points on the moment curve $(t, t^2, t^3, \ldots, t^d)$.*

In particular, for $i = 0, \ldots, \lfloor d/2 \rfloor$,

$$f_i(P) \leq f_i(\Delta(n, d)) = \binom{n}{i+1}.$$

Let us now consider the *edge graph* of a polytope, formed by the polytope's vertices and edges. Theorem 1.1.2 implies in particular that the complete graph K_n, which has all possible edges, is the edge graph of some d-polytope for any $d \geq 4$. This means that these graphs can be arbitrarily dense, in any space of dimension at least four. We can also look at the *dual graph*, where facets are vertices, and two facets are connected by an edge if they share a ridge.

This idea of duality is very important in polytope theory. Suppose

$$P = \mathrm{conv}(\{p_1, \dots, p_n\}) \subseteq \mathbb{R}^d,$$

is a full dimensional polytope containing the origin in its interior. Then the *dual* of P is the polytope D defined as

$$D = \{x \in \mathbb{R}^d : p_i^T x \leq 1 \quad \forall i \in [n]\}.$$

Since every polytope can be described in these conditions via some projection and translation, this dual is well defined. The dual polytope has the same face structure as the *primal* P but inverted: k-dimensional faces become k-codimensional faces and inclusion is reversed. The dual polytope can be interpreted as the set of points corresponding to inequalities that are satisfied within all of P. This has an interesting implication: we may see a polytope as its set of points, or as the set of linear inequalities that it satisfies.

For our purposes we are most interested in *simple* polytopes and *simplicial* polytopes. A polytope is simple if all vertices have degree d in the edge graph. The dual definition is *simplicial* polytopes, where all facets are simplices. A *simplex* (plural, *simplices*) is a convex hull of k affinely independent points, and generalize triangles and tetrahedra.

We could say that simplicial polytopes are the "least singular" polytopes, since the convex hull of points chosen independently at random, sampled from a distribution with dense support over $\mathbb{R}^d$ is simplicial. This is because having $d+1$ points lying in the same hyperplane is an event that happens with probability zero. Also this implies that any polytope is arbitrarily close to a simplicial polytope via small perturbations of the vertices. These perturbations will re-triangulate the existing faces of the original polytope.

The properties of the edge graph of simple polytopes (equivalently, the dual graph of simplicial polytopes) are particularly interesting for the study of an important class of optimization problems, linear programs.

1.1.3 Linear programming and the simplex method

One of the main motivations for the study of polytopes is linear programming.

Definition 1.1.3 (Linear program (LP))**.** *A* linear program *is an optimization problem where the feasible region is a polyhedron and the objective function is linear. This is typically expressed in the* canonical form *as*

$$\begin{array}{cc} \max & c^T x \\ s.t. & Ax \leq b. \end{array}$$

Where $A \in \mathbb{R}^{n \times d}$, $c \in \mathbb{R}^d$ and $b \in \mathbb{R}^n$ are inputs and $x \in \mathbb{R}^d$ are the variables to solve for.

By definition, the optima of such a problem are achieved in a face of the feasible region. Linear programs (LP) are one of the basic classes of optimization problems, and many other algorithms rely on efficient LP algorithms.

Linear programs are known to be solvable in polynomial time on the bit size of the input. This assumes a *bit model* of computation where arithmetic operations are solvable in polynomial time with suitable algorithms.

However, it is still an open problem whether or not they can be solved in *strongly polynomial time*. This is known as the *Smale's 9th problem*. An algorithm runs in strongly polynomial time if it satisfies two conditions: first, it has to run in polynomial time in the *arithmetic model of computation*, where each binary arithmetic operation takes constant time, and the size of inputs and outputs is measured as the number of integer entries it has. Secondly, the bit size of every intermediate result has to be bounded by a polynomial in the bit size of the input. These conditions together automatically imply that the algorithm also runs in *weakly polynomial time* (often called polynomial time), via replacing each arithmetic operation with suitable bit algorithms.

Example 1.1.4. *To clarify this distinction between weak and strong polynomial time, we consider two algorithms that do not run in strongly polynomial time. Each of them fails to fulfill one of the two conditions.*

We start with the Euclidean algorithm for computing the greatest common divisor of n positive integers. It works by successively taking the largest number in the list, then replacing it by its reminder when divided by the smallest number. If this remainder is zero, we discard it from the list, and repeat until only one number is left, the greatest common divisor of the inputs.

The sizes of the intermediate computations are smaller than the original sizes, hence they are bounded by a polynomial, the identity. However, in the arithmetic model of computation it is unknown how many iterations it will take just by looking at the number of inputs. For example, just with two inputs, two consecutive Fibonacci numbers F_{k+1} and F_k, the number of integer divisions computed is k, so we cannot bound this number of divisions just by looking at the number of inputs (2).

But it does run in polynomial time in the bit model of computations. Each integer division reduces the number of bits by at least 1, as the largest number is at least being halved with each iteration. And each such division runs in polynomial time in the total size input. So the number of bit operations performed is polynomial.

The other example is Gaussian elimination for matrices with rational entries. Some variants reduce the intermediate results efficiently (see [22]) and are strongly polynomial. However, a naive implementations lead to numerators and denominators growing exponentially. There is a detailed study of this phenomenon in [24].

Then, this elementary form of Gaussian elimination over the rationals therefore does not run in (weakly) polynomial time either, because it fails at the second condition. Just storing these intermediate results in memory would not be possible in polynomial time. It is worth noting that Gaussian elimination over a finite field runs in polynomial time, as the bit size of intermediate results is bounded by definition and it takes a cubic number of operations.

For linear programming, there are a number of known algorithms running in weakly polynomial time. The first of them was the *ellipsoid method*, but its running time is not very practical for real applications. The best class of algorithms in practice is called in general the *simplex method*.

The simplex method works by exploring the edge graph of the feasible polyhedron, moving from vertex to vertex until it reaches a provably optimum vertex. Recall that since the optima of the linear program form a face, it is guaranteed that some vertex attains the optimum value, at least if the feasible region is bounded and not empty. The particular rules for choosing the next vertex to explore depend on its *pivot rules*, that decide which of the facets in the current vertex will be discarded, and which new facet will be considered. This corresponds to the process of moving through an edge. We call the simplex method a "method" instead of an "algorithm" to recognize that the choice of a different pivot rule changes the complexity of the resulting algorithm.

The simplex method could be a candidate for a strong polynomial algorithm since its number of iterations depends uniquely on the combinatorics of the input polytope. As a matter of fact, it is a strongly polynomial time algorithm if and only if the number of pivot steps is polynomial in the number of variables and constraints of the problem. However, it is not known even if it could run in weakly polynomial time.

Many pivot rules have been suggested, but none of them have been proven to require a polynomial number of pivoting steps. For many of them, pathological polytopes have been found that force the algorithm to visit *all* vertices.

The edge graph itself may have an exponentially large number of vertices as the size of the matrices grow, but the *diameter* of the graph may or may not be polynomial on it. Here the diameter of a graph is the maximum number of steps required to move from any vertex to any other vertex. If this diameter were not polynomial, it would mean that no pivot rule would solve linear programs in a polynomial number of iterations.

1.1.4 Simplicial complexes, simplicial spheres

The topological generalization of simplicial polytopes are *simplicial complexes*. A $(d-1)$-simplicial complex with n vertices is a family of subsets of $[n] = \{1, \ldots, n\}$, called the faces, such that the maximum cardinality of a face is d and subsets of a face are also faces of the same complex. In combinatorics this is often called a *down-set*. A simplicial complex is *pure* if all the maximal faces (the *facets*) have the same dimension, $d-1$ (that is, they have d vertices). Observe that the boundary of a simplicial d-polytope is a $(d-1)$-simplicial complex. We call them $(d-1)$-simplicial complexes to keep the notation consistent with the polytopal case.

Faces of a simplicial complex have special names determined by their dimension, which coincide with the names for polyhedra. We call *Hasse diagram of* $\mathcal{C}$ the Hasse diagram of the partial order of faces by inclusion. That is, it is a directed graph with an arc $f_1 \to f_2$ for every pair of faces f_1 and f_2 with $f_2 = f_1 \cup \{v\}$ for some $v \in V$. Observe that the 1-skeleton and the adjacency graph of a pure complex contain the information about the lower two and the higher two levels of the Hasse diagram, respectively.

A *subcomplex* of $\mathcal{C}$ is a subset of faces that is itself a complex. The subcomplex *induced* by a subset W of vertices is the set of faces of $\mathcal{C}$ contained in W. Associated to every face $f \in \mathcal{C}$ there are the following three (perhaps non-induced) subcomplexes of $\mathcal{C}$.

- The *deletion* of f in $\mathcal{C}$ is the set of faces disjoint from f. That is, it equals the subcomplex induced by $V \setminus f$.
- The *star* of f in $\mathcal{C}$ is the set of faces f' such that $f \cup f'$ is also a face. Equivalently, it is the simplicial complex whose facets are the facets of $\mathcal{C}$ containing f. Observe

that with our definition the star is a *closed* neighborhood of f.

- The *link* of f in $\mathcal{C}$ is the set of faces in the star that do not contain f. That is, it equals the deletion of f in the star of f.

We call *neighborhood* of f the set of vertices of its star.

Every simplicial complex $\mathcal{C}$ has a realization as a topological space obtained as follows: consider a disjoint family of simplices in $\mathbb{R}^N$ (for sufficiently big N) consisting of a simplex of dimension i for each i-face in $\mathcal{C}$. Then take the topological quotient of this set of simplices, by identifying faces as indicated by containment in $\mathcal{C}$. A *simplicial* $(d-1)$-*manifold* or *triangulated manifold* (with or without boundary) is a pure $(d-1)$-complex whose realization is a manifold.

Every ridge of a simplicial manifold is contained in either two or one facets. Ridges of the first type are called *interior* and those of the second type are called *boundary*. The boundary of a $(d-1)$-manifold $\mathcal{C}$ is the pure simplicial $(d-2)$-complex having the boundary ridges of $\mathcal{C}$ as facets. That is, all faces contained in boundary ridges are *boundary faces*. The *adjacency graph* of a simplicial manifold $\mathcal{C}$ is the graph having as nodes the facets of $\mathcal{C}$ and as edges the pairs of facets that share an interior ridge.

A *simplicial* $(d-1)$*-sphere* or $(d-1)$*-ball* is defined in the same way. A simplicial $(d-1)$-sphere is *polytopal* if it can be realized as the boundary complex of a d-polytope. Not all simplicial spheres are polytopal, in fact, most of them are not [41].

A necessary condition for a simplicial sphere to be polytopal is being *shellable*:

Definition 1.1.5. *A* shelling *of a simplicial sphere is an ordering of the facets,* $\mathcal{S} = \{F_1, \ldots, F_n\}$ *such that the complex induced by the first* k *facets,* $k \in \{1, \ldots, n-1\}$ *is always homeomorphic to the* $(d-1)$*-ball. A simplicial complex admitting a shelling is said to be* shellable.

Simplicial polytopes are always shellable: it suffices to look at their dual, and sorting the vertices (facets of the original polytope) by some linear functional. However, not all simplicial spheres are shellable.

1.2 Diameters

Let C be a polyhedron. The *edge graph* of C is the graph defined by the vertices and edges of C. The combinatorial diameter of this graph is the maximum distance, measured in edges traversed between vertices. Some examples for d-polytopes with n facets would be:

Polytope	d	n	Diameter
Polygons	2	n	$\lfloor n/2 \rfloor$
Dodecahedron	3	12	4
d-simplex	d	$d+1$	1
d-cube	d	$2d$	d
d-crosspolytope	d	2^d	2
d-tropical ball	d	$d(d-1)$	$d+2$

In 1957, Warren M. Hirch posed a question to George Dantzig: is the diameter of a d-polyhedron with n facets at most $n-d$? In 1967, Klee and Walkup disproved this for the unbounded case [46]. The same question for polytopes came to be known as the Hirsch conjecture, which remained open until Francisco Santos disproved it in 2012 [63].

As we mentioned in subsection 1.1.3, the complexity of the simplex method is related to the polynomiality of this diameter. It remains unknown if this diameter is bounded by some polynomial in n and d, or even if it is superlineal in $n-d$. The first question is the *polynomial Hirsch conjecture*.

In this section we introduce the known results for various types of complexes, following the material in [45] and [64]. These will prepare the reader for our contributions in chapters 3 and 4.

1.2.1 Reducing to $n = 2d$

Originally we are interested in simple polytopes, since the maximum diameter of polytopes must be attained in a simple one (perturbations of the faces make it simple, and the diameter can only grow). Let $H(n,d)$ be the maximum diameter of (simple) d-polyhedra with n facets. The first key observation is that:

Lemma 1.2.1 (Klee-Walkup,[46]). *$H(n,d) \le H(2(n-d),(n-d))$, with equality if (but not only if) $n < 2d$.*

To prove this we introduce the notions of the *wedge* and *one point suspension* of polytopes:

Definition 1.2.2. *Let $P \subseteq \mathbb{R}^n$ be a polytope with n facets. The* wedge *of P is a $(d+1)$-polytope with $n+1$ facets defined by collapsing a facet F of the product $P \times [0,1]$ (a prism of base P).*

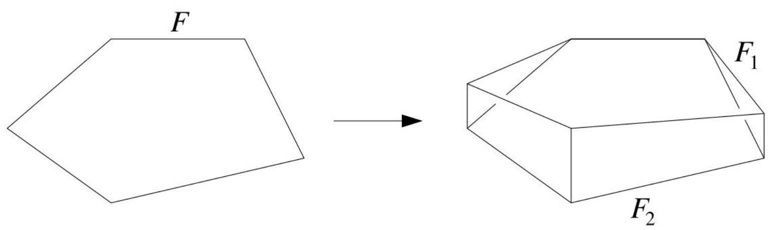

Figure 1.1: Wedge of a polytope over the facet F. Image taken from [45] with permission.

Formally, if F is a facet of P defined by $a^T x \le b$, then the wedge of P over F is:

$$W_F(P) = \{(x,t) \in P \times [0,\infty) : a^T x + t \le b\}.$$

The dual operation of the wedge is the one point suspension*. If P is the convex hull of $p_1, \dots, p_m$, then the one point suspension of P over p_m is defined by embedding P in an affine d-linear space $H \subseteq \mathbb{R}^{d+1}$, and duplicating the vertex p_m, by perturbing it in two opposite directions, perpendicular to H.*

Proof of Lemma 1.2.1. We proceed by induction. Suppose that the d-polytope P has diameter l and n facets.

The wedge of P has at least the same diameter as P, and the dimension and number of facets increased by one. This proves $H(n,d) \leq H(n+1,d+1)$. By iterating this $(n-d)$ times, we conclude that $H(n,d) \leq H(2(n-d),(n-d)$ if $n \geq 2d$.

For $n < 2d$, two vertices of P that achieve a diameter must lie in some common facet F, since each vertex is contained in at least d facets and there is less than $2d$ of them. Looking at this facet, its diameter must be at least the diameter of P, since paths on F are also paths in P. F is then a $(d-1)$-polytope with at most $n-1$ facets and the same diameter as P.

This implies that if $n < 2d$, $H(n,d) \leq H(n-1,d-1)$. Combined with the other direction and induction this gives that if $n < 2d$, $H(n,d) \leq H(2(n-d),(n-d))$.

□

The consequence of this theorem is that we need only to study polytopes with $n = 2d$.

1.2.2 The d-step conjecture and the non-revisiting conjecture

In this section we introduce two (former) conjectures related to the Hirsch conjecture. We will show that they are all equivalent, so the counterexample to the Hirsch conjecture automatically refutes them too.

(Former) conjecture 1.2.3 (d-step conjecture)**.** *Let P be a simple d-polytope with $2d$ facets, and u, v two vertices of P not belonging to the same facet. Then, there is a path of length d in the edge graph connecting them.*

It is clear that the Hirsch conjecture implies this, since such a path cannot have length larger than d (each edge traversed replaces an old facet with a new one, and they do not belong in any common facet), and the Hirsch conjecture gives the upper bound of $n - d = 2d - d = d$.

The other direction also holds by an inductive argument. If the Hirsch conjecture were false (which is the case), then we know thanks to Lemma 1.2.1 that there is a simple d-polytope with $2d$ facets and diameter $> d$. If the two vertices u, v, achieving the diameter are not in common facets, then it would be a counterexample to the d-step conjecture too. If this is not the case, then they must lie in k common facets.

The face of the P in the intersection of these is a simple $(d-k)$-polytope, F. The facets of F correspond to facets of P either containing u or v but not both. Since we know that P is simple and u, v are in k common facets, this gives that F has $2d - 2k$ facets. Therefore, u, v are vertices of a polytope F in the conditions of the d-step conjecture, but with no path of length d between them.

(Former) conjecture 1.2.4 (non-revisiting conjecture)**.** *Let P be a simple d-polytope. Then, for every pair of vertices u, v of P, there is a path in the edge graph between them visits at each step a new facet that has not been visited before.*

This trivially implies the Hirsch conjecture, as the original vertex is in d facets, and with each edge traversed we add a new facet, for a total of at most $n - d$ length.

The converse also holds: let P be a simple d-polytope contradicting the non-revisiting conjecture. We can repeat the argument of Lemma 1.2.1 to assume that it has $2d$ facets. Then, P has two vertices, u, v such that every path between them repeats some facet.

Like in the proof of equivalence for Conjecture 1.2.3, assume that u, v lie in k common facets, and consider the face of P in the intersection of such facets, F. F is a $d-k$ polytope with $2(d-k)$ facets, and the paths between u and v in P have length strictly greater than $d-k$, since it is a counterexample to the non-revisiting conjecture. Then, every path in F between them is also longer than $d-k$, which makes it a counterexample to the d-step conjecture.

The following theorem sums up the results of this section:

Theorem 1.2.5 (Klee-Walkup [46])**.** *The Hirsch, non-revisiting, and d-step conjectures are equivalent.*

1.2.3 Prismatoids and the strong d-step theorem

The d-step theorem of Klee and Walkup (Lemma 1.2.1) reduces the study of the Hirsch conjecture or the asymptotic behavior of $H_{\text{poly}}(n, d)$ to the case $n = 2d$. Santos' construction of non-Hirsch polytopes is based on a version of this result for a particular class of polytopes, the prismatoids.

Note that $H_{\text{poly}}(n, d)$ stands for the edge diameter of simple polytopes, but it also represents the diameter in the *dual graph* of simplicial polytopes. This dual graph is the graph where vertices are facets of the complex, and two vertices are connected if they share $d-1$ vertices (i.e. a ridge). The language of dual diameters of simplicial complexes is combinatorially easier to work with, and so we will use it from now on.

Definition 1.2.6. *A* prismatoid *is a polytope Q with two parallel facets Q^+ and Q^-, that we call the* bases*, containing all the vertices. We call a prismatoid* simplicial *if all faces except perhaps Q^+ and Q^- are simplices. Observe that the faces of a prismatoid of dimension d, excluding the two bases, form a simplicial complex of dimension $d-1$ and homeomorphic to the product of $\mathbb{S}^{d-2}$ with a segment. We call this complex the* prismatoid complex *of Q.*

The width *of a prismatoid is the distance from base to base in the dual graph.*

Theorem 1.2.7 (Strong d-step theorem for prismatoids [63])**.** *If Q is a simplicial d-prismatoid of width l and with $n > 2d$ vertices, there exists a simplicial $n-d$-prismatoid Q' with $2n-2d$ vertices and width at least $l+n-2d$.*

In particular, if $l > d$ then (the simple polytope dual to) Q' violates the Hirsch conjecture.

Proof. This proof follows the proof of Lemma 1.2.1. If both bases are simplicial, then Q already has $n = 2d$, and $l \geq l+n-2d$ trivially. If it's not the case, it must have $n > 2d$. We show how to construct a $(d+1)$-prismatoid with $n+1$ vertices and width at least $l+1$. This construction takes two steps: A one point suspension for one base, and a random perturbation for the other.

If one of the bases is not simplicial, we embed the prismatoid in some d-affine hyperplane in the $d+1$ affine space. Assume that Q^+ is not simplicial, but Q^- may or may not be. Now, we pick a vertex $v \in Q^-$ and duplicate it, such that its two copies v_1, v_2 lie close enough to the original but not in H, and the new base Q'^- remains in a d-dimensional space. Recall that this is the one point suspension of Q^-, dual of the wedge.

For Q^+, we will perturb all vertices while remaining parallel to Q'^-. After the one point suspension of Q^-, but before the perturbation of Q^+, the facets of this intermediate polyhedron are of the following types:

- The two "bases", Q^+ and Q'^-.
- Facets containing both v_1 and v_2, corresponding to facets containing v in the original prismatoid
- Facets containing either v_1 and v_2, corresponding to the facets of the original prismatoid not containing v.

In this intermediate polyhedron, Q^+ is a ridge, since it is contained in a $d-1$ affine space. Then, two of the facets of the third type will be triangulated after we perturb it.

Any path in the new prismatoid Q' connecting the bases projects naturally to a path in the original one, via forgetting one of the v_i. Observe that that there is a final facet required in the higher dimensional path corresponding to a $d-1$-simplex in the triangulated original base Q^+ joined with one of the suspended vertices. However, this simplex just projects to Q^+ after forgetting the v_i. This gives the extra step in the final move to Q'^+. $\square$

Santos' original counterexample [63] applies this result to a 5-prismatoid with 48 vertices and of width six, thus obtaining a non-Hirsch 43-polytope with 86 facets. This was improved in [55] to a 5-prismatoid of the same width but with only 25 vertices, which provides a non-Hirsch polytope in dimension 20.

1.2.4 Upper bounds and normal complexes

Even if our main interest is in the diameter of polytopes, looking at more general classes of pure simplicial complexes is also of interest. We consider now $H_{\text{normal}}(n, d)$, which stands for the maximum diameter of *pure normal simplicial* $(d-1)$*-complexes* with n vertices. A normal simplicial complex is characterized by the property that the link of every face is strongly connected, that is, its dual graph is connected. This property is sometimes called *locally strongly connected* [31]. Recall that d is not the dimension of the complex, but $d-1$ is. One may think of the d as the number of vertices in a facet.

It turns out the two best upper bounds on $H_{poly}(n, d)$, the maximum dual diameter of simplicial d-polytopes with n vertices (which we sometimes denote by $H(n, d)$), work also for normal simplicial complexes. We present these bounds in this section. A simplicial complex is *normal* if the stars of faces are strongly connected, this is, the dual graph of those subcomplexes is connected. Observe that simplicial polytopes (and also simplicial spheres, see Section 1.2.5) are normal trivially, since the star of a vertex is equivalent to its vertex figure, which is also a simplicial polytope.

For this class, we know that:

Theorem 1.2.8 (Kalai-Kleitman [42])**.**

$$H_{poly}(n,d) \leq H_{normal}(n,d) \leq n^{\log_2(d)+1} - 1.$$

Theorem 1.2.9 (Larman [49])**.**

$$H_{poly}(n,d) \leq H_{normal}(n,d) \leq (n-1)2^{d-1}.$$

It is also worth noting than for the class of *flag* normal complexes, the Hirsch conjecture holds. A simplicial complex is *flag* if every clique in the edge graph spans a face.

Theorem 1.2.10 (Adiprasito and Benedetti [1])**.** *Any flag and normal simplicial complex satisfies the non-revisiting path conjecture and, in particular, it satisfies the Hirsch conjecture.*

1.2.5 Simplicial spheres and pure simplicial complexes

Here we talk about diameters for two classes of complexes that will be relevant for chapters 3 and 4.

The first natural generalization of the question on diameters is the case for the most general complexes, the *pure* simplicial complexes. Of course, we have to consider only the case of *strongly connected* pure simplicial complexes, this is, complexes with connected dual graph. For these, it was first known that

$$H_{\text{simp}}(n,d) \leq \frac{1}{d-1}\binom{n}{d-1} - 1.$$

Where $H_{\text{simp}}(n,d)$ is the maximum diameter of pure simplicial $(d-1)$-complexes. Recall that the d here represents the number of vertices in a facet, but the complex is $(d-1)$-dimensional.

This upper bound was asymptotically attained for fixed d in [16]. However, the gap between the bounds grew exponentially in d. In Chapter 3, we show how a probabilistic construction matches more closely the upper bound, to a factor of $\Theta(d^2)$. We also discuss the case for the class of pseudomanifolds, a slightly stronger class of complexes.

A *simplicial* $(d-1)$*-sphere* is a simplicial complex homeomorphic to the $(d-1)$-sphere. Here we are introducing a notion of topology in the complex, by assigning to each facet a geometric simplex and identifying together the common subfaces. If the topology is that of a manifold, we call the complex a *simplicial* $(d-1)$*-manifold*. Simplicial spheres are the natural generalization of the topology of the boundary complex of polytopes. We denote the maximum (dual) diameter among n-vertex $(d-1)$-simplicial spheres as $H_{\text{sph}}(n,d)$.

It is possible to adapt Lemma 1.2.1 to this topological context:

Proposition 1.2.11 (Sphere version of Lemma 1.2.1)**.**

$$H_{sph}(n,d) \leq H_{sph}(2n-2d, n-d) \qquad \forall n,d.$$

Proof. Let $\mathcal{S}$ be a $(d-1)$-sphere with n vertices and let F and G be two facets in it. If $n = 2d$ then there is nothing to prove.

If $n < 2d$ we use induction on $2d - n$. Observe that F and G cannot be disjoint, so let $v \in F \cap G$. Denote $\mathcal{S}' = \mathrm{link}_{\mathcal{S}}(v)$, and consider $F' = F \setminus \{v\}$ and $G' = G \setminus \{v\}$, which are facets in $\mathcal{S}'$. The distance from F to G in $\mathcal{S}$ is clearly bounded by the distance from F' to G' in $\mathcal{S}'$, and $\mathcal{S}'$ is a $(d-2)$-sphere with at most $n-1$ vertices. By inductive hypothesis the diameter of $\mathcal{S}'$ is bounded by $H_{\mathrm{sph}}(2(n-1)-2(d-1),(n-1)-(d-1)) = H_{\mathrm{sph}}(2n-2d, n-d)$.

If $n > 2d$ we use induction on $n - 2d$. Let $\mathcal{S}'$ be the one point suspension of $\mathcal{S}$, whose facets are:

$$\begin{aligned}\mathcal{S}' := \{&F, F \cup \{w_1\}, F \cup \{w_2\}, F \cup \{w_1, w_2\} : F \in \mathrm{link}_{\mathcal{S}}(w)\} \\ \cup \{&F, F \cup \{w_1\}, F \cup \{w_2\} : w \notin F \in \mathcal{S}\}.\end{aligned}$$

Then, $\mathcal{S}'$ is one dimension larger than $\mathcal{S}$, has one more vertex, and at least the same (dual) diameter as $\mathcal{S}$. Therefore,

$$\mathrm{diam}(\mathcal{S}) \le \mathrm{diam}(\mathcal{S}') \le H_{\mathrm{sph}}(n+1, d+1) \le H_{\mathrm{sph}}(2n-2d, n-d),$$

by inductive hypothesis. □

As an immediate consequence,

Corollary 1.2.12. *The Hirsch, d-step, and non-revisiting conjectures for simplicial spheres are equivalent.*

The class of simplicial d-spheres is strictly between normal simplicial complexes and polytopes; recall that most simplicial spheres with enough vertices are not the boundary of any polyhedron [41]. It is known since 1980 that there are simplicial d spheres without the non-revisiting property, hence, topological counterexamples to the Hirsch conjecture:

Theorem 1.2.13 (Mani-Walkup [53])**.** *There is a simplicial* 3*-sphere with* 16 *vertices and without the non-revisiting property. This implies there is a simplicial* 11*-sphere with* 24 *vertices and dual diameter at least* 13*, by the wedge construction.*

In Chapter 4, we present several smaller, lower dimensional non-Hirsch spheres and how we constructed them with an heuristic search. While the Mani-Walkup sphere is known to not be polytopal, the same is not known yet for our examples. But we know they are *shellable*, which is necessary for polytopal simplicial complexes.

More precisely, Table 1.1 sums up the results in this section.

The lower bounds in the last rows of the table are meant asymptotically: they hold for fixed but sufficiently large d and as n goes to infinity. They follow from the known smallest known non-Hirsch polytope ($n = 40$, $d = 20$ [55]) and sphere ($n = 24$, $d = 12$ [53]). We explain how to extend a small non-Hirsch sphere or polytope into a family of spheres (or polytopes) with diameter growing linearly in the next section.

	Lower bound	Upper bound
$H_{\text{simp}}(n,d)$	$\Omega\left(\frac{n^{d-1}}{d^2 d!}\right)$ (Criado-Newman 2019[15], Chapter 3)	$O\left(\frac{n^{d-1}}{d!}\right)$ (Santos 2013 [64])
$H_{sph}(n,d)$	$\simeq 1.08(n-d)$ (Mani-Walkup 1980 [53])	$\min\{2^{d-3}n, n^{\log d-2}\}$ (Larman 1970 [49], Kalai-Kleitman 1992 [42])
	$\simeq 1.11(n-d)$ (Criado-Santos 2019) [17], Chapter 4)	
$H_{poly}(n,d)$	$\simeq 1.05(n-d)$ (Matschke-Weibel-Santos 2017 [55])	

Table 1.1: Known bounds for maximum diameters of classes of simplicial complexes

1.2.6 A linear lower bound

The *excess* of a non-Hirsch simplicial $(d-1)$-sphere (respectively, d-polytope) with n facets and diameter l is defined to be $\frac{l}{n-d} - 1$.

From any non-Hirsch sphere infinitely many additional ones can be obtained by the following procedures:

- The *join* of a (d_1-1)-sphere $\mathcal{S}_1$ with n_1 vertices and diameter l_1 and a (d_2-1)-sphere $\mathcal{S}_2$ with n_2 vertices and diameter l_2 is a (d_1+d_2-1)-sphere with n_1+n_2 vertices and diameter l_1+l_2. It is denoted $\mathcal{S}_1 * \mathcal{S}_2$. The facets of the join of two complexes are defined in the following way: the vertices of each facet is the union of the vertex sets of two facets, one from each complex. In a polytopal sense, the join of two polytopes is exactly the dual of the Cartesian product of the two polytopes: facet normals of the join of two polytopes are concatenations of a facet normal from the first side with a facet normal from the second one.
- The *connected sum* of two $(d-1)$-spheres $\mathcal{S}_1$ and $\mathcal{S}_2$ with n_1 and n_2 vertices and of diameters l_1 and l_2 is a $(d-1)$-sphere with n_1+n_2-d vertices and of diameter l_1+l_2-1. It is denoted $\mathcal{S}_1 \# \mathcal{S}_2$. (Strictly speaking, the diameter of a connected sum of two spheres depends on the choice of facets to glue; we here assume that they are glued in the worst possible way). This operation is always realizable polytopally by deforming both polytopes via a projective transformation such that the faces to be glued are big enough that the remaining vertices are close to them, thus guaranteeing polytopality after the identification.
- The *two points suspension* of $\mathcal{S}$ has one more dimension and two more vertices than $\mathcal{S}$, and its diameter is one more than that of $\mathcal{S}$. It is defined by embedding the polytope in a higher dimensional space and building a bipyramid over it.

These constructions, which all preserve polytopality, lead to the following:

Theorem 1.2.14 (Variation of [63, Theorem 6.5]). *If for a certain d_0 we know that $H_{sph}(2d_0, d_0) = l_0$, then*

$$H_{sph}(n,d) > \left\lfloor \frac{n-d}{d} \right\rfloor \cdot \left(\left\lfloor \frac{d}{d_0} \right\rfloor (l_0 - d_0) + d - 1 \right) \simeq \frac{l_0}{d_0}(n-d), \qquad \forall n, d.$$

In particular, we have $H_{sph}(n,d) \gtrsim (n-d)\frac{l_0}{d_0}$ *for* $n \gg d \gg d_0$.

The same holds for H_{poly}.

Proof. Let $\mathcal{S}_0$ be the initial sphere, of dimension $d_0 - 1$, with $2d_0$ vertices, and with diameter l_0. Then, for every k the k-fold join $\mathcal{S}^{*k} = \mathcal{S} * \cdots * \mathcal{S}$ *(k times)* of $\mathcal{S}$ has dimension $kd_0 - 1$, diameter kl_0, and $2kd_0$ vertices. For a given d, letting $k = \lfloor d/d_0 \rfloor$ and performing $d - d_0 k$ two point suspensions on $\mathcal{S}^{*k}$ produces:

$$H_{\mathrm{sph}}(2d,d) \geq \left\lfloor \frac{d}{d_0} \right\rfloor (l_0 - d_0) + d, \qquad \forall d,k.$$

Let T^d be the $(d-1)$-sphere on $2d$ vertices obtained so far. By a connected sum of $\lfloor n/d \rfloor - 1 = \lfloor (n-d)/d \rfloor$ copies of T^d we conclude that

$$\begin{aligned} H_{\mathrm{sph}}(n,d) \geq H_{\mathrm{sph}}(d\lfloor n/d \rfloor, d) &\geq \left\lfloor \frac{n-d}{d} \right\rfloor \cdot \left(\left\lfloor \frac{d}{d_0} \right\rfloor (l_0 - d_0) + d \right) - \left\lfloor \frac{n-d}{d} \right\rfloor + 1 \\ &= \left\lfloor \frac{n-d}{d} \right\rfloor \cdot \left(\left\lfloor \frac{d}{d_0} \right\rfloor (l_0 - d_0) + d - 1 \right) + 1. \end{aligned}$$

And the same arguments apply to polytopes. □

1.3 Tropical geometry

In this section we introduce the elements of tropical geometry that we will use especially in Chapter 2. We refer the reader to [35] for a more complete introduction of tropical geometry.

1.3.1 The tropical semiring, tropical varieties and tropical linear spaces

The *tropical semiring* is the triplet $\mathbb{T} = \{\mathbb{R} \cup \{\infty\}, \oplus, \odot\}$ where $a \odot b = a + b$ and $a \oplus b = \min(a,b)$. We may consider $\oplus$ to be max or min, so it is customary in this field to always declare which one it represents. We prefer the *min convention.* In the *max convention*, the second operator is max and the zero of the semiring is $-\infty$ instead.

This satisfies all ring axioms except existence of inverses for $\oplus$. ∞ becomes the zero element, as $\infty \odot x = \infty$ and $\infty \oplus x = x$, for all $x \in \mathbb{T}$. 0 is the unit since $0 \odot x = x$ for all $x \in \mathbb{T}$. Multiplicative inverses for any non-∞ values exist and are unique (since $x \odot (-x) = 0$), so the semiring is sometimes called a *semifield.* Also observe that tropical addition is idempotent, and tropical exponentiation corresponds to a scalar multiplication: $a^{\odot n} = na$.

The motivation for this language is twofold. First, it provides an algebraic structure for optimization and graph problems involving linear functions with integer coefficients. The second reason comes from algebraic geometry. The tropical semiring is the natural algebraic structure to study the image of *valued fields.*

A *valued field* is a field K equipped with a *valuation function* $v : K \to \mathbb{T}$ such that:

- $v(a) = \infty$ if and only if $a = 0$,
- $v(ab) = v(a) + v(b) = v(a) \odot v(b)$,
- $v(a + b) \geq \min(v(a), v(b)) = v(a) \oplus v(b)$, with equality if $v(a) \neq v(b)$.

The most interesting example of a valued field is the field of Puiseux series $K\{\{t\}\}$ (power series of the form $\sum_{k=0}^{\infty} c_k x^{k/n}$), where the valuation is the degree of the smallest monomial. This valuation map provides discrete and computational techniques for the study of algebraic geometry.

For the purposes of this thesis, we focus on the first motivation.

With addition and multiplication defined, we can extend naturally to tropical linear algebra. The *tropical semimodule* is the set of tuples $\mathbb{T}^d$ equipped with element-wise addition and tropical scalar product:

$$(a_1, \ldots, a_d) \oplus (b_1, \ldots, b_d) = (a_1 \oplus b_1, \ldots, a_d \oplus b_d),$$

$$\lambda \odot (a_1, \ldots a_d) = (\lambda \odot a_1, \ldots \lambda \odot a_d).$$

We can also define the dot product of two vectors as

$$\langle (a_1, \ldots, a_d), (b_1, \ldots, b_d) \rangle_{\odot} = \oplus_{i=1}^{d} (a_i \odot b_i).$$

Matrix multiplications are defined equivalently, by taking the dot products of columns and rows. It will also be of interest later on to define the power of a matrix as $A^{\odot k} = A \odot \cdots \odot A (k \text{ times})$

Tropical polynomials are defined analogously. A *d-variate tropical polynomial* is an expression of the form

$$F(x) = \oplus_{u \in S} (a_u \odot x_1^{\odot u_1} \odot \cdots \odot x_d^{\odot u_d}).$$

Where $S \subseteq \mathbb{Z}^d$ is the set of exponents. In other notation:

$$F(x) = \min_{u \in S} (a_u + u_1 x_1 + \cdots + u_d x_d).$$

The *support* $\text{supp}(F)$ of a tropical polynomial is the subset of S with nonzero (i.e. not ∞) coefficients. We may assume that this support is never the empty set, that is, F takes real values only.

The tropical analogue of the vanishing set of a polynomial is defined as follows. A tropical d-variate polynomial F is said to *vanish* at a point $p \in \mathbb{R}^d$ if the minimum is attained twice, this is, two of the tropical monomials achieve that minimum. The *tropical hyperpersurface* defined by F is the set of points that make it vanish.

Observe that since F is piecewise linear, the tropical hypersurface corresponds to the $(d-1)$-skeleton of the *dome* of F (this is, the graph of F seen as a polyhedral complex). This implies that tropical hypersurfaces are polyhedral complexes. Observe also that in the definition of tropical polynomial we require that no two monomials have the same exponents, so no cell will be full dimensional.

We are particularly interested in the tropical hyperplanes. They are defined by linear polynomials, this is, expressions of the form

$$F(x) = \min_{i \in d}(a_i + x_i).$$

Analogously to the classical projective geometry, we may quotient $\mathbb{R}^{d+1}$ (or $\mathbb{T}^{d+1}$) by the equivalence relation defined by multiplication by a scalar. This defines the *tropical projective torus* $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. This means considering tuples of $d+1$ real numbers just looking at their relative differences, so $a \simeq b$ if and only if $a + \lambda\mathbb{1} = b$ for some real λ. One way to find representatives for these classes is by considering the classical hyperplane of the points with sum of coordinates 0. This gives us a way to represent graphically the varieties of tropical (homogeneous) polynomials. The other classical way to choose representatives is by setting a fixed coordinate to 0.

Tropical hyperplanes are not homeomorphic to classical hyperplanes. In the tropical projective d-torus, they look like the union of the $(d-1)$-faces of the face fan of a regular d-simplex, and the *apex* of the fan may be any point in the tropical projective torus.

1.3.2 Tropical polytopes and polytropes

We proceed now to define the tropical analogues for convexity and halfspaces. For a set of points $p_1, \ldots, p_n \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$, a *tropical convex combination* of them is a point of the form $\oplus_{i \in [n]}(\lambda_i \odot p_i)$ for some $\lambda \in \mathbb{R}^n$. Observe that unlike with classical combinations, we do not require the coefficients, λ to be positive or add up to one. A set closed under tropical linear convex combinations is called *tropically convex*. Note also that tropically convex sets may not be classically convex, nor classically convex sets are necessarily tropically convex.

The tropical analogue of a polytope is then the class of *tropical polytopes*, which are tropical convex hulls of finitely many points.

Another class of tropically convex sets are the (closed) tropical halfspaces [34]. They admit multiple equivalent definitions; we prefer to define them as the set $\{x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} : a^T \odot x \leq b^T \odot x\}$ for two vectors $a, b \in \mathbb{R}^{d+1}$. Observe that for any given index $i \in [d+1]$, we may assume without loss of generality that either $a_i = \infty$ or $b_i = \infty$. Note also that the locus of points satisfying the expression with equality is a subset of the tropical hyperplane defined by $(a \odot b)^T x$. In other words, we may see them like a union of maximal cones of the face fan of the regular simplex. An *open tropical halfspace* is topologically the interior of the closed version, and it is described by the same inequality in the strict form.

Then, it becomes natural to define the *tropical polyhedra*, which are intersections of finitely many tropical halfspaces. We may describe them as sets of the form $\{x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} : A\dot{x} \leq B \odot x\}$ where $A, B \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}^{nx(d+1)}$. Much like in the classical case, these correspond exactly with tropical polytopes. One direction is trivial, as tropical half spaces are tropically convex (a property preserved by intersections). The other direction is harder to prove, but it also holds. These two statements can be summed up as follows:

Theorem 1.3.1 (Tropical Minkowski-Weyl theorem [35])**.** *The intersection of finitely many tropical halfspaces in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ is a tropical polytope. Conversely, any tropical polytope admits such a representation*

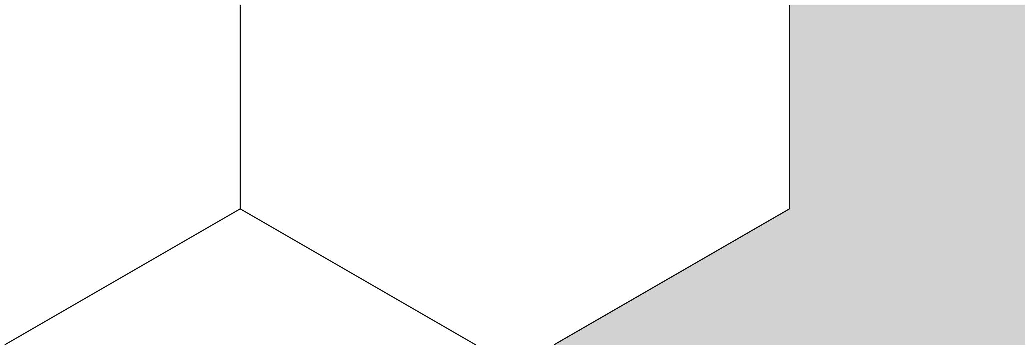

Figure 1.2: Left: a tropical linear space subdivides the tropical projective torus into $(d+1)$ regions. Right: A tropical halfspace is a union of these regions.

A set of points that is at the same time a classical and tropical polytope is called a *polytrope*. These are particularly interesting for our interests. Observe that we can discard tropical polytopes defined by half spaces of the tropically non-convex type, as such half spaces must be touching the polytope in a ridge only. Then, a polytrope has at most $d(d-1)$ classical facets, with normal vectors of the form $e_i - e_j$, differences of vectors of the canonical base. The description as tropical polytope may have less facets, as we can group together faces with the same i under the same tropical half space.

Then we have a compact way of describing these polytropes, with a matrix $A \in \mathbb{T}^{d \times d}$ with zeroes in the diagonal. The polytrope is then the set of points satisfying $p_i - p_j \leq A_{ij}$ for all $i, j \in [d+1]$. Or, using linear algebra notation, $p^T \odot A \odot (-p) \geq 0$, which is remarkably similar to the classical equation for ellipsoids.

1.3.3 Certifying empty polytropes

The next natural question for polytropes is to determine if they are empty or not. While this is equivalent to linear programming for general polyhedra, in this case the normal vectors are restricted so we may exploit this extra structure.

We will use a variant of Farkas' lemma to certify emptiness:

Theorem 1.3.2 (Farkas' lemma (variant))**.** *Let $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$. Exactly one of these systems of inequalities has solutions:*

- $Ax \leq b$, $x \in \mathbb{R}^d$,
- $A^T y = 0$, $b^T y < 0$, $y \leq 0$, $y \in \mathbb{R}^n$.

In this particular case, the matrix A is already determined; its rows are of the form $e_i - e_j$ for e_i, e_j vectors of the canonical base of $\mathbb{R}^{d+1}$. We assume that $A_{i,i} = 0$. A solution of the second system, which would be a certificate of emptiness of the polytrope, assigns a value $y_{i,j}$ to each pair $i \in [d+1], j \in [d+1], i \neq j$ in such a way that $A^T y = 0$, this is,

$$\sum_{j'=0}^{d+1} y_{i,j'} - \sum_{i'=0}^{d+1} y_{i',i} = 0.$$

We may look at these values as a *circulation* over the complete (oriented) graph K_{d+1} with $d+1$ vertices. A circulation assigns values to each oriented edge in such a way the total "flow" getting in and out of each vertex is zero. The most basic type of circulation would be a cycle of the graph, with all values of y being the same in that cycle.

It turns out that every circulation on a graph can be factored as a sum of cycles. An easy way to prove this is recursively. Assume the minimum flow the circulation assigns is attained in some edge $e \in E(K_{d+1})$. We start looking at the origin of e, then we move towards its destination. Now, for every vertex we visit, we move in the direction of any edge with non zero flow. Since e had the minimum flow among all the edges in the graph, they always carry at least as much flow, and since the sum of flows in each node is zero, some outwards edge must exist. The number of vertices in K_{d+1} is limited so eventually we return to some vertex we have visited already. This closes a loop. If we subtract to the ys the value of the minimum flow in the cycle, the number of edges with non zero flow just decreased by at least one. We repeat this process and eventually reach a decomposition of the circulation as a sum of cycles.

This means that in order to see if a polytrope is empty, it just suffices to look for cycles in K_{d+1} with negative sum of edges in the polytrope matrix A. Thus, the next problem to solve is how to find these cycles efficiently.

Observe that the tropical matrix multiplication $A^{\odot 2} = A \odot A$ has in each entry $A^{\odot 2}_{i,j}$ the length of the shortest path from j to i, because the tropical dot product corresponds to taking the minimum among all possible intermediate nodes l, of the paths from j to l and to i. We can iterate this observation and note that $A^{\odot k}$ has the length of the shortest paths from every vertex to any other vertex in exactly k steps. We define the *all-pairs shortest path matrix* of A as

$$A^* = A \oplus A^{\odot 2} \oplus A^{\odot 3} \oplus \dots .$$

Which encodes the shortest paths in any number of steps. A matrix of this form, is called a *shortest path matrix.* If A has no negative cycles, then this expression converges, since it is enough to consider paths of at most $(d+1)$ edges. Any longer path will have a cycle, which must have positive cost. Conversely, if we compute $\oplus_{k=1}^{d+1} A^{\odot k}$, and it has no negative element in the main diagonal, this result must be A^* for the same reasons.

An efficient way to compute it is by iterating the operator $X \mapsto X \oplus (X \odot X)$, $\lceil \log(d+1) \rceil$ times to the matrix A. This takes $O(\log(d)d^3)$ time. It is worth noting that there is a simple and more efficient $O(d^3)$ algorithm using dynamic programming, the Floyd-Warshall algorithm.

This computation also gives us the values of all the linear programs with objective functions of the form $e_i - e_j$. Geometrically it means that the matrix A^* is the "closest" of all the possible representations of the same polytope, as each classical halfspace in that description is tangent to it.

1.3.4 Tropical linear programming

A *tropical linear program* is an optimization problem of the form:

$$\begin{array}{ll} \min & c^T \odot x \\ \text{s.t.} & A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-. \end{array}$$

With $A^+, A^- \in \mathbb{T}^{n\times(d+1)}$, $b^+, b^- \in \mathbb{T}^n$, $c \in \mathbb{T}^{d+1}$. It is possible to adapt the simplex method to the tropical setting [4], and then finding an optimum point corresponds to exploring a simplicial sphere by traversing its dual graph. This is another reason to motivate the study of the diameter of other classes of simplicial complexes (and in particular, simplicial spheres) in chapters 3 and 4.

It must be said however, that tropical polytopes do not have exactly the same combinatorics as classical convex polytopes. The boundary of a tropical polytope may not be a pure simplicial complex, and the notion of face is not defined yet. But the vertices and edges of the tropical polytrope are well defined. And, in classical geometry a simple polytope is uniquely determined by its graph [10]. In this sense, they do have a similar combinatorial structure induced by their graphs.

Like in the classical case, these problems are equivalent to a satisfiability problem, in this case, the *tropical linear satisfiability problem.* The latter problem asks for a solution to a list of tropical linear inequalities or a certificate that such solution does not exist. This is the exact formulation of this second problem:

$$\exists ? x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} : A \odot x \geq B \odot x.$$

The tropical linear satisfiability problem arises naturally from mean-payoff games [2], and it has applications in the field of scheduling problems. This problem is known to be in both NP and co-NP [2], and no polynomial time algorithm is known for it yet.

We may consider without loss of generality that A has exactly one non-∞ term in each row. If this were not the case, we could split that inequality in two, since the tropical dot product is computing a minimum of these two. In the same way we can assume the corresponding element in B to be ∞. If an element is not ∞ in a row of A, and the corresponding element of B is not ∞ either, then we can either ignore the entry in B (if the number in B is smaller than the one in A) or the row is itself an unsatisfiable inequality. Then, the inequalities become classified in $d+1$ types, by the index in A in which the row has its non-∞ entry. Every row of B will have at least one ∞ entry, the one corresponding to its type, and at least one non-∞ entry (otherwise the row would be a trivial constraint).

Furthermore, we can assume also that there are rows of every type. If one type were missing, we can consider the corresponding coordinate to be arbitrarily small, which will make every inequality easier to satisfy. In practice this produces another, lower-dimensional, tropical linear feasibility problem.

A geometric approach to the problem is based on the *Shapley operator.* It is defined as the map $S : \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \to \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ such that:

$$S(x) = -B^T \odot^{\min} (A \odot^{\max} x).$$

Intuitively speaking, the Shapley operator projects the point in each of the $d+1$ cardinal directions until it satisfies all inequalities of that type, and one of them with

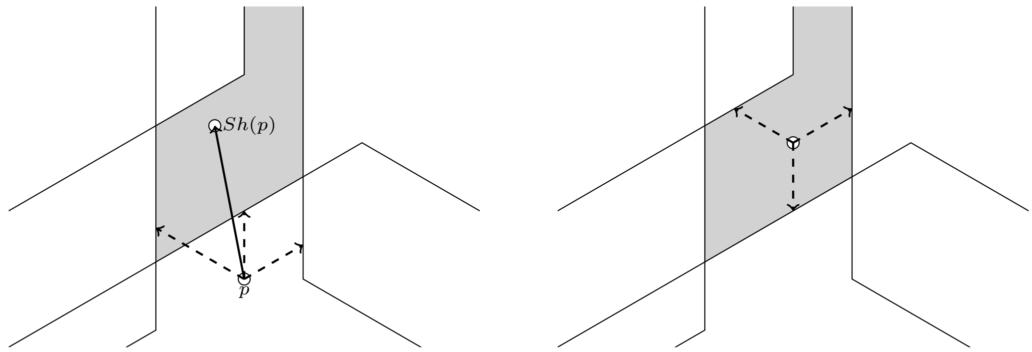

Figure 1.3: Left: a tropical feasibility problem with constraints of all types, and the Shapley operator of a point p. Right: A fixed point of the Shapley operator in this problem is equidistant to the 3 types of constraints

equality. Then, the next point is the current one, plus the $d + 1$ differences between the current point and the projections.

A fixed point of this operator is either a positive or negative certificate, because it violates every type of inequality by the same amount. If the amount is positive, then the point violates at least one inequality of each type. These together cover up the whole space. And if the amount is negative, then obviously all inequalities are satisfied. The existence of a fixed point can be proven using topological arguments and continuity of the Shapley operator.

Just iterating this operator enough times makes the point orbit around a fixed point. This gives an algorithm, follow a trajectory for long enough and average the points it inspects.

When observed under the lens of the *tropical distance*, $\text{dist}(a, b) = \max_{i\in[d+1]}(a_i - b_i) - \min_{j\in[d+1]}(a_i - b_i)$ it gains additional meaning. The distance between the images of two points under the Shapley operator is at most that of the distance between the original points, by a routine calculation. By applying this to the fixed point (which we know exists), then the Shapley operator is guaranteed to give a point *at least as close* to the fixed point as the original.

This suggest a separation oracle: Either the point is already a fixed point, or it gives us a closer one. If we study the *tropical bisector* separating them, it could give us information on which points to consider next. In classical linear feasibility the separating object is a classical hyperplane, which is at the same time a bisector, but this is no longer the case in tropical geometry. Tropical bisectors are then the natural way to adapt convex optimization algorithms to the tropical world.

1.4 Computational geometry

Computational geometry is the branch of geometry and algorithm design that studies algorithms and data structures for geometric objects. It has obvious applications in computer assisted design (CAD) and optimization.

In this section we talk about the algorithmic foundations that we will use for our results, and some known data structures that inspired our algorithms for tropical Voronoi diagram computations. In particular we introduce the *Voronoi diagram*, which is a central object in this thesis. Our main reference for elementary computational geometry is *Computational Geometry: Algorithms and Applications*[19].

As is customary in this field, we use the *arithmetic model* of computation, where each arithmetic operation takes a constant amount of time, independently of the precision of their operands or their type.

1.4.1 Convex hulls

We refer to the definitions in section 1.1 for convex hull, and polytope. Even if we have discussed the basic properties for a convex hull of a set $S \subset \mathbb{R}^d$ of points, we still have not said how to compute them. We first introduce a simple algorithm for the 2-dimensional case:

Algorithm 1.4.1 (Andrew's algorithm)**.** *Let $S \subseteq \mathbb{R}^2$ be a set of n points in the plane. Our goal is to compute $C = (C_1, \dots, C_k)$ be the ordered list of vertices in the convex hull in counter-clockwise order. Hence, $C \subseteq S$. Let us assume that $S = \{p_1, \dots, p_n\}$ is sorted by some generic linear functional. Without loss of generality, this functional is $f(x) = x_1$.*

We compute the 'upper' part of the hull first, this is, the faces whose normal vector form a positive angle with $(1, 0)$. Or alternatively, the normal vector points 'upwards'. The 'lower hull' is analogous.

First, we initialise $C = p_1, p_2$, this is, the two leftmost points. Now, for each point in S, in order from left to right, we append it to the hull. If the angle formed by the last three points is positive, we move on to the next point. If the angle is negative, we remove the point where the negative angle is formed (the second to last point) and repeat this procedure until the angle is positive.

Eventually we reach the last point p_n. Since each point has been inserted in the list exactly one time (and removed at most one time), the running time of this march is $O(n)$, plus a sorting time of $O(n \log n)$ for sorting the points.

This procedure is repeated for the lower hull, but traversing the list in reverse order.

The running time of $O(n \log n)$ is the best we can achieve. This is by reduction to sorting (which is known to have at worst that complexity). Assume we want to sort the list of real numbers $\{x_1, \dots, x_n\}$. If we compute the convex hull of $\{(x_1, x_1^2), \dots, (x_n, x_n^2)\}$, and we traverse it in order starting at the lowest point, we retrieve the sorted list.

For higher dimensions, convex hull computations become complex. This is because the complexity of describing the face complex grows exponentially in the dimension (as well as the number of facets). It is still an open problem to find a convex hull algorithm that runs in output-sensitive polynomial time.

This example of the convex hull illustrates also a common theme in computational geometry: *what happens if the angle formed by three consecutive points is not positive or negative, but zero?*. The problem is that there is more than one combinatorial way to describe the same geometric object if the input points are not given in *general position*.

A common way to deal with these degeneracies is via *symbolic perturbations*, adding symbolic errors to each input number that will help break ties in each comparison. Since these perturbations have been added before running the algorithm, the tie breaking is consistent, so the combinatorial structure of the output is as expected. Another possibility in this case is to consider only the *minimal* convex hull.

Since dealing with these degenerate cases may be a non-trivial problem, we often assume the input to be in *general position*, free of these degeneracies. What this means depends on the problem at hand, and sometimes on the algorithm that we are using. Weaker generality conditions are better (since they means that the algorithm runs successfully in more cases). This generality condition must be general enough that all possible inputs are arbitrarily close to a general position one. This is, the set of general position inputs is dense in the space of possible inputs, for some topology meaningful for the problem at hand.

1.4.2 Voronoi diagrams

Another classical object in computational geometry is the Voronoi diagram. The Voronoi diagram of a set of *sites* $S = \{p_1, \ldots, p_n\}$ in a metric space (for the time being, $\mathbb{R}^d$ with Euclidean distance) is the partition of space into n regions such that points in region k are closer or equidistant to p_k than to any other p_i.

It is easy to see the interest of such diagrams. They serve to inspect local and metric properties of spaces, and for point location problems. They are also relevant in clustering problems, and they appear in nature, in particular when the sites of each region are allowed to move to the centroid of the corresponding region. In some applications, other metrics besides the classical Euclidean distance are used, like max-distance and Manhattan distance, or arbitrary polytopal metrics.

For now, let us inspect the geometric properties of the Euclidean case. First of all, each region is a polyhedron, because for each pair of sites p_i, p_j, the condition "being closer (or equidistant) to p_i than to p_j" defines a linear halfspace (since bisectors in the Euclidean distance are linear spaces). The region of p_i is the intersection of these halfspaces, which is a polyhedron. Observe that these polyhedra may not be bounded (and indeed, the regions corresponding to sites in the boundary of the convex hull of all sites are unbounded).

No two Voronoi regions will have full dimensional intersection, since the distance to two distinct points is a different function, hence there is a direction where perturbing the point increases one distance and decreases the other. And faces of each region correspond to loci of the points that are equidistant to certain subsets of S. This induces a polyhedral complex structure in the diagram, where facets are the polyhedral regions and lower dimensional faces are their intersections. Each face is uniquely determined by the subset of sites it is closest to.

If we assume that no $d+2$ sites are in the same sphere (i.e. for every $d+2$ sites, no point of $\mathbb{R}^d$ is equidistant to all of them), then the complex is simple. In this case, the dual complex is the *Delone triangulation* (or *Delaunay triangulation*) of S.

Vertices in the Voronoi diagram of S translate to d-simplices in the corresponding Delone triangulation. The Delone triangulation of S is then the unique triangulation of S formed by triangles whose circumscribed circles do not contain any site of S in their

Figure 1.4: A Voronoi curiosity: the current territorial subdivision of Spain was designed by Javier de Burgos in 1822 [11] with two objectives: each province is the Voronoi region of its respective capital, and each capital is the centroid of its corresponding province, both measured with travelling time. Nowadays, this subdivision has barely changed, and it is one of the oldest national territorial subdivisions in the world. The two conditions are commonly seen whenever Voronoi diagrams appear in nature.

interior. One way to see that it is unique and a triangulation comes from the duality with Voronoi diagrams. A more formal way to prove it uses Theorem 1.4.2.

1.4.3 Fortune's algorithm

We present here a *sweep line algorithm* for the 2-dimensional Delone triangulation problem. For some direction v chosen in general position, we consider a line with normal vector v that moves in that direction *sweeping* all points in S.

The abstract idea is to keep track of the locus of points that have equal distance to the sweep line and to the visited points of S, i.e. the subset of S that the sweep line has already gone through. We call this locus the *beach line*. This beach line is a monotone curve formed by segments of parabolas, with foci in the sites, and the sweep line as their directrix. It is then represented as the ordered list of the foci of these parabolas. Notice that the same point may appear more than once.

We maintain a priority queue of events that may happen as we advance the sweep line. These events are of two types: A parabola arc may disappear between two other arcs, or a new parabola may be inserted as the sweep line crosses a new site.

Since the endpoints of adjacent parabolas arcs trace the Voronoi Diagram, the first type of event adds a new vertex in the diagram. The second type of event adds two parabola arcs to the sweep line: The one corresponding to the inserted vertex and the arc that was just split in two.

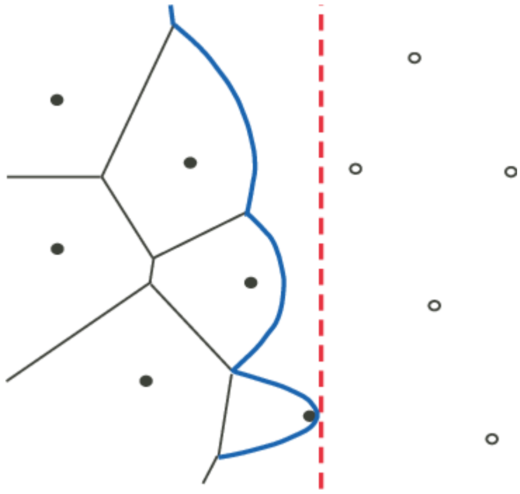

Figure 1.5: Fortune's algorithm mid-execution. The dotted line is the *sweep line* and the blue line is the *beach line*

The size of the beach line is never bigger than $2n$, as each insertion adds two arcs, and each deletion removes one. Looking for the place where a new arc is to be inserted takes time $O(\log n)$ by binary search. The priority queue registers two new inserted events with every swept site, and just an insertion with every removed arc. This means that there is a linear number of events considered, since the number of vertices in the Voronoi diagram grows linearly with n.

This gives a total time complexity of $O(n \log n)$ time, as the initial sorting takes $O(n \log n)$, the priority queue grows at most linearly, and the processing of an event takes at most $O(\log n)$.

There are other Voronoi diagram algorithms in two dimensions with the same worst time complexity as this one (divide and conquer, randomized, or reducing it to a convex hull as we will see in the next section). The advantage of Fortune's algorithm is that it generalizes well for other metrics, specially if they are strictly convex, this is, spheres in these metrics contain no segments. If the metric isn't strictly convex, some singularities may occur unless we ask for stronger regularity conditions.

1.4.4 Connection to convex hulls

The two introductory computational geometry problems we have have talked about, the convex hull and the Voronoi diagram (equivalently, Delone triangulations) are closely tied by the following result:

Theorem 1.4.2. *Let $S \subseteq \mathbb{R}^d$ be a finite set of points in general position (i.e. no $d+2$ points are in the same sphere, and they affinely span $\mathbb{R}^d$). Consider the set $S^+ = \{(x_1, \dots, x_d, ||x||^2) : x \in S\}$. Then the triangulation induced by the lower facets of* $\mathrm{conv}(S^+)$ *is exactly the Delone triangulation of S.*

Proof. We will prove how a simplex in the lifted convex hull corresponds to a simplex in the Delone triangulation, this is, simplex whose circumscribed sphere contains no points of S in its interior.

Let us look at $d+1$ sites of S. First of all we will prove that we may assume them to lie on the unit sphere $\mathbb{S}^{d-1}$. We can move the sites by subtracting a vector $c \in \mathbb{R}^d$ and a scaling $r \in \mathbb{R}$, the circumcenter and circumradius of the simplex. Now we study how such a transformation affects the last coordinate in S^+:

$$\left\| \left(\begin{array}{c|c} I/r & 0 \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c|c} I & -c \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c} x \\ \hline 1 \end{array}\right) \right\|,$$

$$= \left\langle \left(\begin{array}{c|c} I/r & 0 \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c|c} I & -c \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c} x \\ \hline 1 \end{array}\right), \left(\begin{array}{c|c} I/r & 0 \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c|c} I & -c \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c} x \\ \hline 1 \end{array}\right) \right\rangle,$$

$$= \left\langle \left(\begin{array}{c|c} I/r & -c/r \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c} x \\ \hline 1 \end{array}\right), \left(\begin{array}{c|c} I/r & -c/r \\ \hline 0 & 1 \end{array}\right) \left(\begin{array}{c} x \\ \hline 1 \end{array}\right) \right\rangle = \frac{1}{r^2}\langle x, x\rangle - \frac{2c}{r^2}x + 1.$$

The last coordinate in S^+ also changed affinely, plus this affine transformation in $d+1$ dimensions has positive determinant (Its matrix is triangular). This means that the transformation did not change the orientations of simplices, and if $d+1$ sites of S^+ form a facet of the convex hull, they also form it after this affine transformation.

Thus we assume without loss of generality that the $(d+1)$-tuple of points lie in the unit sphere of $\mathbb{R}^d$. Now the unit sphere corresponds exactly to the circumsphere of these points, and the last coordinate of the raised sites is exactly the distance to the circumcenter squared. In order for these points to be a facet, this last coordinate has to be larger than 1, since the sites in the tuple have height 1. This corresponds exactly to being farther away than 1 unit to their circumcenter, which completes the proof. □

The previous proof has various immediate consequences: First, it shows that the Delone triangulation of S in general position is unique. Second, it gives an immediate algorithm for computing Delone triangulations (and Voronoi diagrams) in any dimension, provided we know how to compute convex hulls of arbitrary dimensions.

Also, since the combinatorial complexity of the convex hull is known (by the upper bound theorem, Theorem 1.1.2) it also tells us that the number of faces of all dimensions in the Voronoi diagram grows like $O(n^{\lceil d/2 \rceil})$ for fixed d.

1.4.5 Randomized incremental constructions

Consider the following algorithmic problem: Given a set of n points in $\mathbb{R}^2$, what is the smallest circle that contains them all?

For points in *general position* (here meaning that no four points form a cyclic quadrilateral), such circle is defined by either 2 points defining its diameter, or 3 points.

If we are *guaranteed* that two particular points, say, p_1 and p_2 are in the boundary of the optimal circle, a linear time algorithm is easy to describe. Just initialise the circle as the one with diameter $p_1 p_2$ and iterate through all the remaining points. If the next point is inside the current smallest circle, we ignore it. Otherwise, we update it as the circumscribed circle of the p_1, p_2 plus the new point. Since p_1 and p_2 are in the boundary by assumption, this circle has to contain all previously examined points.

What if we are *guaranteed* that a particular point, p_1, is in the boundary? We can think analogously. First, we set the circle to have diameter p_1p_2 for any p_2, and iterate through all points. If a point is inside the current circle, we ignore. If it is not, then it is *guaranteed* that the new point, p_k is in the boundary of the smallest circle containing $p_1, \ldots, p_k$. If it were not, p_k would be in the interior of the smallest circle containing $p_1, \ldots, p_{k-1}$, since each point is either in the interior or in the boundary of such circle.

Apparently this algorithm takes $O(n^2)$ time to run. But if we shuffle the points randomly before starting, the expected running time becomes linear. This is because each newly inserted point p_k has a probability of at most $2/(k-1)$ of triggering the first algorithm. So, in expected time, each iteration takes constant time and globally the algorithm still runs in linear time.

The same idea generalizes for the problem without any information on border points. We start with the diameter of the first two points. Every time we add a point, p_k, it is either inside the circle (and we ignore it) or outside. If it is outside, which happens with probability at most $3/k$ then we *know* that it is in the border of the smallest circle covering $p_1, \ldots, p_k$, and we call the corresponding algorithm that runs in linear time.

These key ideas of randomization of input, incrementally adding points, and having a low expected number of expensive updates are what inspired the *randomized incremental algorithms*. These algorithms are commonly used in various problems in computational geometry, especially when working with geometric data structures.

Intuitively this approach works because adding data in a randomized incremental way gives rough global approximations, and further data only refine the local details.

This approach has been successfully used for convex hull computations as well. First, randomly rearrange the points, then compute the convex hull of the first $d+1$ points and iteratively pick up the next point and examine its relative position with respect to the constructed polytope so far. If it is in the right side of all facets, we ignore it. If it is in the wrong side of some facet, we retriangulate this facet with the new point, then flip out all concave faces until the result is a convex polytope again.

The complexity analysis of this randomized incremental convex hull is difficult in arbitrary dimension, but in three dimensions it becomes easy: The amount of work done by inserting a vertex is proportional to the degree of that vertex in the convex hull after insertion, and also counting for the amortized cost of potentially deleting it later on. The average degree in the edge graph of a simplicial 3-polytope with n vertices is approximately 6, by Euler's formula and simpliciality.

Then, the most expensive part of each iterations is finding a facet that this vertex is violating. The number of facets is linear in the number of vertices (again, by Euler's formula), so the complexity of this convex hull algorithm is $O(n^2)$. With appropriate data structures for locating the facet, it is $O(n \log n)$.

We will use the randomized incremental approach to describe efficient algorithms for tropical Voronoi diagram computations in Chapter 2.

Chapter 2

Tropical bisectors and Voronoi diagrams

This chapter is from the preprint "Tropical bisectors and Voronoi diagrams" [14], by Francisco Criado, Michael Joswig and Francisco Santos. It is available at `https://arxiv.org/abs/1906.10950`.

2.1 Introduction

One early route to the success of tropical geometry is based on the tropicalization of classical algebraic varieties defined over some valued field. Key examples include Mikhalkin's correspondence principle, which relates tropical plane curves with classical complex algebraic curves [56], or the tropical Grassmannians of Speyer and Sturmfels [66]. In all of this the focus lies on the combinatorial properties of tropical varieties, which are ordinary polyhedral complexes.

More recently, however, tropical semi-algebraic sets and their intrinsic geometry came into the picture; cf. [3], [33]. For instance, their metric properties appear in [5] as a tool to show that standard versions of the interior point method of linear programming exhibit an exponential complexity in the unit cost model. The proof of this result is based on translating metric data on a family of tropical linear programs into curvature information about the central paths of their associated ordinary linear programs. Similarly, tropical analogs of isoperimetric (or isodiametric) inequalities have been studied in [20], where a tropical volume is defined that corresponds to an "energy gap" in mathematical physics [47]. Another example is the statistical analysis of phylogenetic trees by Lin, Monod and Yoshida [50].

We feel that all this calls for a more systematic investigation of metric properties of tropical varieties. Starting from first principles, this naturally leads to tropical Voronoi diagrams.

The *tropical distance* between two points $a, b \in \mathbb{R}^{d+1}$ is

$$\operatorname{dist}(a,b) \;=\; \max_{i\in[d+1]} (a_i - b_i) - \min_{j\in[d+1]} (a_j - b_j) \;=\; \max_{i,j\in[d+1]} (a_i - b_i - a_j + b_j) \;. \quad (2.1)$$

It does not depend on choosing min or max as the tropical addition. The map dist : $\mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \to \mathbb{R}$ is non-negative, symmetric, and it satisfies the triangle inequality. Moreover, it is homogeneous, so it induces a norm on the *tropical d-torus* $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \cong \mathbb{R}^d$, where $\mathbb{1} = (1, \dots, 1)$ denotes the all ones vector. The *tropical Voronoi region* of a site $s \in S$ with respect to a set S comprises those points in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ to which s is the nearest among all sites in S, with respect to dist. The *tropical Voronoi diagram* $\mathrm{Vor}(S)$ is the cell decomposition of $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ into Voronoi regions. Tropical Voronoi diagrams are a special case of Voronoi diagrams for polyhedral norms, a classical topic in convexity and computational geometry; cf. [7, Sect. 7.2] or [54, Sect. 4].

The intersection of two or more Voronoi regions is part of a *bisector*, i.e., the locus of points which are equidistant to a given set. For instance, in the Euclidean case the bisector of two points is a degenerate quadric which agrees with an affine hyperplane as a set. In the tropical setting, the bisector of two points can also be described as part of a tropical hypersurface, but this is now of degree $d + 1$; cf. Proposition 2.4.1. Further, in the tropical case two points may already produce degenerate bisectors (which may contain, e.g., full-dimensional pieces), whereas the first degenerate case in the Euclidean metric arises for three points. So tropical Voronoi diagrams behave quite differently from Euclidean Voronoi diagrams.

Yet there are also similarities. A key structural result is that the tropical Voronoi regions are star convex and can be described as unions of finitely many ordinary polyhedra; cf. Proposition 2.3.1 and Theorem 2.4.14. We prove a second main result, Theorem 2.3.18, for the more general case of an arbitrary polyhedral norm in $\mathbb{R}^d$: the bisector of any three points in weak general position is homeomorphic to an open subset of $\mathbb{R}^{d-2}$. Our proof generalizes the arguments from [29], [30], where a similar result was proved for smooth norms in $d = 2, 3$. However, the global topology of tropical bisectors of three or more points can be radically different from the topology of the classical bisectors. For instance, tropical bisectors are sometimes disconnected and, more strongly, $d + 1$ points can have more than one circumcenter. This may happen even in general position; cf. Examples 2.3.12 and 2.3.26. We do not know if bisectors may have nontrivial higher Betti numbers, but we suspect they can; cf. Theorem 2.3.24.

Another contribution is a randomized incremental algorithm for computing the tropical Voronoi diagram of n points in general position in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ with an expected runtime of order $O(n^d \log n)$, for fixed dimension d; cf. Theorem 2.5.13. Euclidean Voronoi diagrams of finite point sets can be explained fully in terms of ordinary convex polyhedra and convex hull algorithms; cf. [19], [7]. We do not know if there is a tropical analog.

The chapter is organized as follows. The short Section 2.2, in which we verify that the tropical distance is induced by a polyhedral norm and discuss the combinatorics of the tropical unit ball, sets the stage. In Section 2.3 we collect our general structural results on bisectors and Voronoi diagrams. The results in this section are proved for general polyhedral norms, but all our examples address the tropical case. A subtle point is the right concept of "general position". In fact, we distinguish between *weak general position* which prevents bisectors to contain full-dimensional parts (cf. Proposition 2.3.5), and a (stronger) *general position* which is defined via stability of bisectors under small perturbations of the sites. For instance, the bisector of any number k of points in general position in $\mathbb{R}^d$ is a polyhedral complex of *pure* dimension $d+1-k$; cf. Corollary 2.3.8. As a special case, the bisector of $d + 1$ points in $\mathbb{R}^d$ in general position is finite. Section 2.4 returns to the tropical case. We specialize our results on bisectors in general polyhedral norms, and we show that the combinatorial types of tropical bisectors of two points are classified in terms of a certain polyhedral fan related to the tropical unit ball and the braid arrangement; cf. Theorems 2.4.6 and 2.4.14. This is related to work of Develin [21] on the moduli of tropically collinear points. Finally, in Section 2.5 we discuss algorithms. This includes a tropical variant of Fortune's beach line algorithm [25] for planar Voronoi diagrams as well as the aforementioned algorithm in arbitrary dimension.

2.2 The tropical unit ball

The unit ball with respect to the tropical distance function dist defined in (2.1) is

$$\begin{aligned} \mathbb{B}^d \ &= \ \left\{ x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \;\middle|\; \operatorname{dist}(x,0) = 1 \right\} \ = \ \bigcap_{i \neq j} \left\{ x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \;\middle|\; x_i - x_j \leq 1 \right\} \\ &= \ \operatorname{conv}\left(\{\pm 1\}^{d+1} \setminus \{\pm \mathbb{1}\} \right) + \mathbb{R}\mathbb{1} \ . \end{aligned} \tag{2.2}$$

In this way, $\mathbb{B}^d$ is a polytope in the tropical projective torus $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. We also write $\mathbb{B}^d(a,r)$ for the tropical ball with center a and radius r. All tropical balls result from scaling and translating $\mathbb{B}^d$. In fact, the tropical norm equals the polyhedral norm with respect to the tropical unit ball, in the sense of Section 2.3. Such distances are called *convex distance functions* in [7, Sect. 7.2]; see also [28, 30, 29].

Both the inequality and the vertex descriptions of $\mathbb{B}^d$ in Equation (2.2) are non-redundant:

- $\mathbb{B}^d$ has $d(d+1)$ facets. Each facet corresponds to a choice of coordinates achieving the maximum and the minimum.
- $\mathbb{B}^d$ has $2^{d+1}-2$ vertices. Each vertex corresponds to a (nontrivial) partition of the coordinates into maxima and minima. For example, $\mathbb{B}^2$ is a hexagon and $\mathbb{B}^3$ is a rhombic dodecahedron.

The vertex description also shows that $\mathbb{B}^d$ equals the projection of the $(d+1)$-dimensional regular cube $[-1,1]^{d+1}$ in $\mathbb{R}^{d+1}$ along the direction $\mathbb{1}$. That is, $\mathbb{B}^d$ is a zonotope with $d+1$ generators in general position, and all its faces are parallelepipeds. These generators correspond to the $d+1$ coordinate directions in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. This suggests a combinatorial way to specify the faces of $\mathbb{B}^d$: Each face F can be written as a Minkowski sum

$$F \ = \ \sum_{i=1}^{d+1} s_i \ ,$$

where each s_i is one of $\{-e_i\}$, $[-e_i, e_i]$ or $\{+e_i\}$. We say that F is of *type* (F_-, F_*, F_+) if

$$\begin{aligned} F_- \ &= \ \{ i \in [d+1] : s_i = \{-e_i\} \} \ , \\ F_* \ &= \ \{ i \in [d+1] : s_i = [-e_i, e_i] \} \ , \\ F_+ \ &= \ \{ i \in [d+1] : s_i = \{e_i\} \} \ . \end{aligned} \tag{2.3}$$

Conversely, a partition of $[d+1]$ into three parts F_-, F_*, F_+ corresponds to a face of $\mathbb{B}^d$ if and only if neither F_- nor F_+ is empty. Moreover, the dimension of F equals the cardinality of F_*. In particular, the vertices of $\mathbb{B}^d$ correspond to the $2^{d+1}-2$ ways

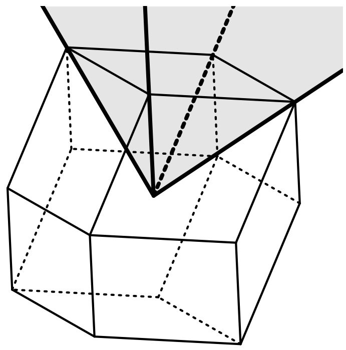

Figure 2.1: The tropical 3-ball $\mathbb{B}^3$, with the conical hull of a facet highlighted

of partitioning $[d+1]$ into two non-empty subsets. The facets of $\mathbb{B}^d$ correspond to the $d(d+1)$ ways of choosing an ordered pair from $[d+1]$, without repetition.

Remark 2.2.1. *The zonotope $\mathbb{B}^d$ is dual to an arrangement of $d+1$ linear hyperplanes in general position in $\mathbb{R}^d$, oriented so that the intersection of all positive half-spaces is empty. In particular, its face lattice is the same as the lattice of non-zero covectors of the unique totally cyclic oriented matroid of rank d with $d+1$ elements. Covectors of an oriented matroid are usually written as (V_-, V_0, V_+) but in our context we prefer to use $*$ instead of zero meaning that the corresponding coordinate is not fixed.*

Remark 2.2.2. *Another general description of $\mathbb{B}^d$ is that it equals the (ordinary) Voronoi cell of the lattice of type A_d (i.e., the triangular lattice for $d = 2$ and the face centered cubic lattice (FCC) for $d = 3$). Similarly, $\mathbb{B}^d$ is the polytope polar to the difference body $T - T$ of a regular d-simplex T. This description shows that $\mathbb{B}^d$ is the same as the polytope U_d that appears in Makeev's conjecture. See, e.g., [67, Conjecture 21.3.2].*

2.3 Bisectors in polyhedral norms

Throughout this section we work in the general framework of Minkowski norms; cf. [7, Sect. 7.2], [28], [54]. Consider a convex body $K \subset \mathbb{R}^d$ with the origin in its interior. Let $\mathrm{dist}(a, b)$ be the unique scaling factor $\alpha > 0$ such that $b - a \in \alpha \partial K$. Then, dist satisfies the triangle inequality, is invariant under translation, and homogeneous under scaling. If $K = -K$ then $\mathrm{dist}(a, b) = \mathrm{dist}(b, a)$ and $\mathrm{dist}(0, \cdot)$ is a norm in $\mathbb{R}^d$ in the usual sense. We allow $K \neq -K$, whence $\mathrm{dist}(a, b) \neq \mathrm{dist}(b, a)$, but we still call it a norm. Bisectors and Voronoi diagrams for these norms have been studied in computational geometry [7, Sect. 7.2], [54, Sect. 4].

For any finite point set S we define:

$$\mathrm{bisector}(S) \;:=\; \left\{ x \in \mathbb{R}^d \;\middle|\; \mathrm{dist}(a, x) = \mathrm{dist}(b, x) \text{ for } a, b \in S \right\} \quad .$$

Following the computational geometry tradition we will often call the elements of S the *sites*. Although most of our results do not need this, for simplicity we assume K to be a polytope. In this case we denote by $\mathcal{F}(K)$ the face fan of K. The norm $\operatorname{dist}(0,\cdot)$ is linear in each of these cones, so we write

$$\operatorname{bisector}_{F_1,\dots,F_k}(\{a_1,\dots,a_k\}) \;=\; \operatorname{bisector}(\{a_1,\dots,a_k\}) \cap a_1 + F_1 \cap \dots \cap a_k + F_k \;\;,$$

for the intersection of the bisector with a choice of cones $F_i \in \mathcal{F}(K)$. Each cell of the form $\operatorname{bisector}_{(F_1,\dots,F_k)}(a_1,\dots,a_k)$ is the intersection of the polyhedron $(a_1+F_1)\cap\dots\cap(a_k+F_k)$ with an affine subspace, which implies it is itself a polyhedron. As a consequence:

Proposition 2.3.1. *Let K be a polytope with the origin in its interior, and let* dist *be the corresponding Minkowski norm. Let $S = \{a_1,\dots,a_k\} \in \mathbb{R}^d$ be a finite point set. Then the set* $\operatorname{bisector}(\{a_1,\dots,a_k\})$ *is a polyhedral complex whose cells are the polyhedra*

$$\operatorname{bisector}_{(F_1,\dots,F_k)}(a_1,\dots,a_k)$$

for all choices of $F_1,\dots,F_k \in \mathcal{F}(K)$.

Proof. The family of polyhedra

$$\operatorname{bisector}_{(F_1,\dots,F_k)}(a_1,\dots,a_k)\,,\ \text{ with } F_1,\dots,F_k \in \mathcal{F}(K)\,,$$

forms a polyhedral complex since

$$\operatorname{bisector}_{(F_1,\dots,F_k)}(S) \cap \operatorname{bisector}_{(F_1',\dots,F_k')}(S) \;=\; \operatorname{bisector}_{(F_1\cap F_1',\dots,F_k\cap F_k')}(S) \;\;.$$

That polyhedral complex covers the entire bisector since for each point $p \in \operatorname{bisector}(S)$ and for each i, the point a_i must lie in some face F_i of $p - \operatorname{dist}(a_i,p)K$. $\square$

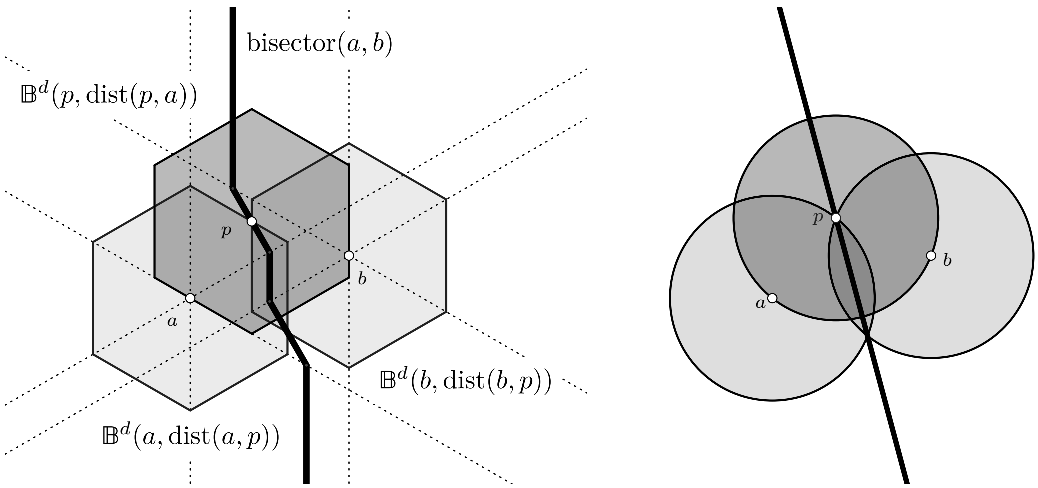

Figure 2.2: Left: A point, p, in the tropical bisector of a and b. Right: The analogous picture, in classical geometry.

Our primary example is the case where $K = \mathbb{B}^d$ is the tropical ball. In Figure 2.2 the point p, which is generic within the bisector of a and b, lies in the facet $\operatorname{bisector}_{(-*+),(+-*)}(a,b)$.

For the purpose of representing $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ in 2-d pictures we must choose some suitable linear map. Any three vectors $v_1, v_2, v_3 \in \mathbb{R}^2$ with $v_1 + v_2 + v_3 = 0$ define a map from $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ to $\mathbb{R}^2$ via $e_i \mapsto v_i$. While, e.g., $v_1 = (1,0)$, $v_2 = (0,1)$, $v_3 = (-1,-1)$ is a common choice in tropical geometry, for our pictures, such as Figure 2.2 (left), we settle for the more symmetric

$$v_1 = \left(-\sin\frac{2\pi}{3}, \cos\frac{2\pi}{3}\right) \ , \ v_2 = \left(\sin\frac{2\pi}{3}, \cos\frac{2\pi}{3}\right) \ , \ v_3 = \big(0,1\big) \ .$$

Remark 2.3.2. *If* $\operatorname{bisector}(a_1, \ldots, a_k) = \emptyset$ *then* $\operatorname{bisector}(a'_1, \ldots, a'_k) = \emptyset$ *for every sufficiently small perturbation of the points.*

2.3.1 Weak general position and general position

Definition 2.3.3 (General position)**.** *A finite point set $S \subset \mathbb{R}^d$ is in* general position *with respect to K, if for every subset $\{a_1, \ldots, a_k\} \subset S$ there are neighborhoods U_i of each a_i such that for every choice of $\{a'_1, \ldots, a'_k\}$ with $a'_i \in U_i$ and for every choice of maximal cones $F_1, \ldots, F_k \in \mathcal{F}(K)$ we have*

$$\operatorname{bisector}_{(F_1,\ldots,F_k)}(a_1, \ldots, a_k) = \emptyset \iff \operatorname{bisector}_{(F_1,\ldots,F_k)}(a'_1, \ldots, a'_k) = \emptyset \ .$$

Moreover, the set S is in weak general position *if no pair of points $a, b \in S$ lies in a hyperplane parallel to a facet of K.*

Remark 2.3.4. *As the name suggests, "weak general position" is implied by "general position" (see Corollary 2.3.6). A yet stronger notion would arise requiring stability not only of facets but also of lower-dimensional cells in bisectors; that is, allowing lower-dimensional cones from $\mathcal{F}(K)$ for the F_i in the definition of general position. But this intermediate notion of "general position" is the most appropriate for our purposes and the algorithms in Section 2.5. As a first indication, Theorem 2.3.7 provides a local characterization.*

By Proposition 2.3.1 bisectors are polyhedral complexes, and thus each cell has a dimension.

Proposition 2.3.5. *Two points a, b are in weak general position if and only if* $\operatorname{bisector}(a, b)$ *does not contain full-dimensional cells.*

Proof. Since $\operatorname{bisector}_{(F,F')}(a, b) \subset (a + F) \cap (b + F')$, for it to be d-dimensional we need F and F' to be cones of facets. We also need $F = F'$, so that $\operatorname{dist}(a, \cdot)$ and $\operatorname{dist}(b, \cdot)$ have the same gradient on $(a + F) \cap (b + F)$, and we need $b - a$ to be parallel to the facet, so that $\operatorname{dist}(a, \cdot) = \operatorname{dist}(b, \cdot)$ on $(a + F) \cap (b + F)$.

Conversely, if $b - a$ is parallel to a facet of K with cone F then

$$\operatorname{bisector}_{(F,F)}(a, b) \ = \ (a + F) \cap (b + F) \ ,$$

and this is d-dimensional. □

Corollary 2.3.6. *General position implies weak general position.*

Proof. Suppose $a - b$ is parallel to a facet of K, so that $\mathrm{bisector}_{(F,F)}(a, b)$ is full-dimensional, where F is the cone of that facet. Taking b' close to b but away from the hyperplane parallel to the facet makes $\mathrm{bisector}_{(F,F)}(a, b')$ empty. Now the claim follows from Proposition 2.3.5. $\square$

Theorem 2.3.7. *Let $S = \{a_1, \ldots, a_k\} \subseteq \mathbb{R}^d$ and for each a_i choose a maximal cone $F_k \in \mathcal{F}(K)$. Let $Q := (a_1 + F_1) \cap \cdots \cap (a_k + F_k)$, let $\lambda_{F_i}(x)$ be the linear function that restricts to $\mathrm{dist}(0, x)$ on F_i, and let H be the affine subspace defined by $\lambda_{F_1}(x - a_1) = \cdots = \lambda_{F_k}(x - a_k)$.*

Then, the following conditions are equivalent:

1. *There are neighborhoods U_i of each a_i such that for any choice of $a'_i \in U_i$ the polyhedron $\mathrm{bisector}_{(F_1,\ldots,F_k)}(a'_1, \ldots, a'_k)$ is not empty.*
2. *(a) Q is full-dimensional and H intersects its interior; and*
   (b) the $k - 1$ functions $\lambda_{F_i} - \lambda_{F_1}$ for $i = 2, \ldots, k$ are linearly independent.

Since
$$\mathrm{bisector}_{(F_1,\ldots,F_k)}(a_1, \ldots, a_k) \;=\; Q \cap H \tag{2.4}$$
condition (a) is equivalent to "$\mathrm{bisector}_{(F_1,\ldots,F_k)}(a_1, \ldots, a_k)$ meets the interior of Q".

Proof. For the implication from "1" to "2", let us first show that "1" forces Q to be full-dimensional. Aiming for a contradiction, we assume that Q is contained in the boundary of one of the cones $a_i + F_i$. That is, the polyhedron $Q_i := \bigcap_{j \neq i}(a_j + F_j)$ does not meet the interior of $a_i + F_i$, for some i. Then any a'_i in the interior of $a_i + F_i$ yields $(a'_i + F_i) \cap Q_i = \emptyset$. Hence $\mathrm{bisector}_{(F_1,\ldots,F_k)}(a'_1, \ldots, a'_k) = \emptyset$, where $a'_j = a_j$ for $j \neq i$. This contradicts "1" and shows that Q is full-dimensional.

To see that $\mathrm{bisector}_{(F_1,\ldots,F_k)}(a_1, \ldots, a_k)$ must intersect the interior of Q, suppose we are given neighborhoods U_i as in "1"; recall (2.4). For each i, choose v_i in the interior of F_i and such that $\lambda_i(v_i) = 1$. Let $a'_i := a_i + \varepsilon v_i$, where $\varepsilon > 0$ is taken small enough so that $a'_i \in U_i$. Observe that our choice of v_i makes the affine subspace defined by $\lambda_{F_1}(x - a'_1) = \cdots = \lambda_{F_k}(x - a'_k)$ agree with H. Let $Q' := (a'_1 + F_1) \cap \cdots \cap (a'_k + F_k)$, which lies in the interior of Q. We have
$$\mathrm{bisector}_{(F_1,\ldots,F_{d+1})}(a'_1, \ldots, a'_{d+1}) \;=\; Q' \cap H \;\; .$$
By "1" this is not empty, and thus H intersects the interior of Q.

For condition (b) observe that H can equivalently be defined by the $k - 1$ affine equalities
$$(\lambda_{F_i} - \lambda_{F_1})(x) = \lambda_{F_i}(a_i) - \lambda_{F_1}(a_1) \qquad \text{for } i = 2, \ldots, k \;\; . \tag{2.5}$$
If the left-hand sides are linearly dependent, then one of the $k - 1$ equations, say the ith one, is redundant. But then choosing a point $a'_i \in U_i$ with $\lambda_{F_i}(a'_i) \neq \lambda_{F_i}(a_i)$ (and letting $a'_j = a_j$ for the rest) renders the system of equations infeasible. Hence $\mathrm{bisector}_{(F_1,\ldots,F_{d+1})}(a'_1, \ldots, a'_{d+1}) = \emptyset$, contradicting "1".

We now show that "2" implies "1". Consider arbitrary points a'_i, and let $Q'_i = \bigcap_i (a'_i + F_i)$. Further let H' be the affine subspace defined by
$$\lambda_{F_1}(x - a'_1) \;=\; \cdots \;=\; \lambda_{F_k}(x - a'_k) \;\; ,$$

so that

$$\text{bisector}_{(F_1,\ldots,F_k)}(a'_1,\ldots,a'_k) \;=\; Q'\cap H' \;.$$

We want to show that if each a'_i is sufficiently close to the corresponding a_i for all i then $Q'\cap H'$ is not empty.

Condition (b) says that H' is $(d+1-k)$-dimensional (and parallel to H) for every choice of a'_is and that it varies continuously with the choice. Condition (a) says that Q' stays full-dimensional if the a'_is are sufficiently close to the original a_is, and that it also varies continuously with the choice, in the following strong sense: consider a description of each $a_i + F_i$ by a finite system of linear inequalities. Then $a'_i + F_i$ is defined by a system with the same linear functions and with right-hand sides varying continuously with the a'_is.

Thus, if each a'_i is close to a_i, then Q' and H' are a full-dimensional polyhedron and a $(d+1-k)$-dimensional affine subspace close to Q and H respectively. Since, by (a), H intersects the interior of Q, continuity implies that H' still intersects the interior of Q' when a'_i is close enough to a_i for each i. In particular, $\text{bisector}_{(F_1,\ldots,F_k)}(a'_1,\ldots,a'_k)$ is not empty. □

Corollary 2.3.8. *The bisector of k points in general position is either empty or pure of dimension $d+1-k$. In particular, the bisector of $d+1$ points in general position is finite, and this is empty for more than $d+1$ points.*

Proof. Every maximal non-empty cell $\text{bisector}_{(F_1,\ldots,F_k)}(a_1,\ldots,a_k)$ is the intersection of the polyhedron Q with the affine subspace H of Theorem 2.3.7. That result implies that H has dimension $d+1-k$ and meets the interior of the full-dimensional polyhedron Q. Thus, the cell has dimension $d+1-k$. □

Corollary 2.3.9. *If every subset of at most $d+2$ points in S is in general position then so is S.*

Corollary 2.3.10. *For any $n \geq 1$, the sequences of n points in $\mathbb{R}^d$ in general position form an open dense subset of $(\mathbb{R}^d)^n$.*

Proof. Let $S \subset \mathbb{R}^d$ be a set of cardinality n. Then S is in general position if and only if, for each subset $\{a_1,\ldots,a_k\} \subset S$ and maximal cones $F_1,\ldots,F_k \in \mathcal{F}(K)$, the polyhedron $\text{bisector}_{(F_1,\ldots,F_k)}(a_1,\ldots,a_k)$ either is empty or satisfies condition (1) of Theorem 2.3.7. Since there are finitely many choices of $\{a_1,\ldots,a_k\}$ and $\{F_1,\ldots,F_k\}$, it suffices to prove the statement for one such choice.

For openness: Emptyness is an open condition because of the continuity of $\text{dist}(\cdot,\cdot)$. On the other hand, Theorem 2.3.7 says that condition (1) is equivalent to (2), which is open by the arguments in the proof (namely, the fact that both Q and H depend continuously on the sites).

For density, we are going to show that if $Q\cap H$ is not empty but fails to satisfy condition (a) or (b) then the set $\{a_1,\ldots,a_k\}$ lies in one of finitely many linear hyperplanes in $(\mathbb{R}^d)^k$. To emphasize that Q and H depend on the choice of sites we denote them $Q(a_1,\ldots,a_k)$ and $H(a_1,\ldots,a_k)$.

If (b) fails for $\{a_1,\ldots,a_k\}$ then the linear system (2.5) defining $H(a_1,\ldots,a_k)$ is feasible but overdetermined, which implies a linear relation, depending solely on $F_1,\ldots,F_k$, among the right-hand sides $\lambda_{F_i}(a_i) - \lambda_{F_1}(a_1)$. The relation is not tautological on the a_is since, as shown in the proof of Theorem 2.3.7, it is easy to construct a point set $\{a'_1, \dots, a'_k\}$ with $H(a'_1, \dots, a'_k) = \emptyset$.

For (a), consider the inequality descriptions of the cones F_i and translate them to obtain an inequality description of $Q(a_1, \dots, a_k)$ as the feasibility region of a system $\mathcal{S}(a_1, \dots, a_k)$ of affine inequalities with fixed gradients and with right-hand sides parameterized linearly by the a_is. If (a) fails for $\{a_1, \dots, a_k\}$ then one of two things happen:

- *$Q(a_1, \dots, a_k)$ is non-empty but not full-dimensional.* Consider a minimal subsystem of $\mathcal{S}(a_1, \dots, a_k)$ that already defines a non-full-dimensional feasibility region. Minimality implies that this feasibility region is an affine subspace of $\mathbb{R}^d$ and that turning the inequalities to equalities produces an over-determined subsystem. This implies, as in the previous case, a linear relation among the a_is.

- *$Q(a_1, \dots, a_k)$ is full-dimensional but $Q(a_1, \dots, a_k) \cap H(a_1, \dots, a_k)$ is contained in its boundary.* This implies that $Q(a_1, \dots, a_k) \cap H(a_1, \dots, a_k)$ is contained in a facet of $Q(a_1, \dots, a_k)$. Let $H_0(a_i)$ be the hyperplane containing that facet. Note that the hyperplane H_0 depends only on one a_i since it comes from the description of one of the cones $a_i + F_i$. Then, adding to the the linear system (2.5) the equation defining $H_0(a_i)$ produces again an over-determined system, hence there is a linear relation among the a_is.

The relations for condition (a) are not tautological on the a_is since we can easily make $Q(a'_1, \dots, a'_k)$ full-dimensional and $H(a'_1, \dots, a'_k)$ intersect its interior as follows: choose each a'_i in the interior of $-F_i$ and such that $\lambda_{F_i}(a_i) = -1$, so that both H and the interior of Q contain the origin. The relations are finitely many since we have at most one for each subsystem (respectively equation) of the system $\mathcal{S}(a_1, \dots, a_k)$, which depends only on the choice of F_is. $\square$

For the case of the tropical norm, condition (b) admits a nice combinatorial characterization. Observe that a choice of facets $F_1, \dots, F_k \in \mathcal{F}(\mathbb{B}^d)$ can be encoded as a directed graph on the vertex set $[d+1]$ and with an arc a_i going from the coordinate that is minimized at F_i to the coordinate that is maximized at F_i, for $i = 1, \dots, k$. We denote this graph $G(F_1, \dots, F_k)$.

Proposition 2.3.11. *For the case of the tropical norm, condition 2.(b) of Theorem 2.3.7 holds if and only if the graph $G(F_1, \dots, F_k)$ either has no (undirected) cycle or it has a unique cycle and it is unbalanced; that is, the number of arcs in one direction is different from the other direction.*

Proof. A cycle in $G(F_1, \dots, F_k)$ is equivalent to a linear dependence among the corresponding linear functions λ_{F_i}s, by simply adding them with signs corresponding to the direction of the arcs along the cycle. If the cycle is balanced then λ_{F_1} can be subtracted from each λ_{F_i} so that the corresponding functions $(\lambda_{F_i} - \lambda_{F_1})$ are also dependent. The same thing can be done if $G(F_1, \dots, F_k)$ has two different (unbalanced) cycles, since a linear combination of the two corresponding dependences can be made balanced.

Conversely, any linear dependence among the functions $(\lambda_{F_i} - \lambda_{F_1})$ corresponds to a balanced dependence among the corresponding λ_{F_i}. The latter either corresponds to a balanced circuit in the graph or decomposes into two (or more) linear dependences with distinct supports. $\square$

As a consequence of Corollary 2.3.8 the bisector of a set S of $d+1$ points in general position is a finite set of points, which we call *circumcenters* of S. In dimension two, three points in (weak) general position have a most one circumcenter, as we show in Corollary 2.3.19 below. In higher dimension the same is known to be false for other polytopal norms [29], and here is an example for the tropical norm:

Example 2.3.12 (Non-uniqueness of circumcenters). *Let us consider the four points $a_1 = (0,2,3,3)$, $a_2 = (0,4,2,2)$, $a_3 = (2,4,1,1)$ and $a_4 = (4,0,2,2)$. Their bisector contains the points $x = (0,0,1,-1)$ and $y = (0,0,-1,1)$. Indeed, both x and y are at distance 4 from all the a_i's since we have*

$$\begin{aligned} a_1 - x &= (0,2,2,4), & a_1 - y &= (0,2,4,2),\\ a_2 - x &= (0,4,1,3), & a_2 - y &= (0,4,3,1),\\ a_3 - x &= (2,4,0,2), & a_3 - y &= (2,4,2,0),\\ a_4 - x &= (4,0,1,3), & a_4 - y &= (4,0,3,1). \end{aligned}$$

The points $a_1, \dots, a_4$ are not in weak general position, as they lie in the plane $x_3 - x_4 = 0$. However, they satisfy conditions (a) and (b) of Theorem 2.3.7 for the polytopes Q_x and Q_y containing the circumcenters x and y: For condition (b) observe that the digraphs corresponding to x and y are, respectively, $\{14, 12, 32, 21\}$ and $\{13, 12, 42, 21\}$. They both have a single cycle, $\{12, 21\}$, which is unbalanced. Condition (a) follows from the fact that x (resp. y) is in the interior of all the cones whose intersection defines Q_x (resp. Q_y). This is equivalent to the fact that all the vectors $a_i - x$ and $a_i - y$ have a unique maximum and a unique minimum entries.

Since condition (b) does not depend on the sites and condition (a) is, by Theorem 2.3.7, open, any perturbation of the sites will still produce at least two circumcenters. In particular, by Corollary 2.3.10, there are sites in general position for which this happens.

2.3.2 Halfspheres, sectors, and the bisector of two points

The topology of a bisector is closely related to the following partition of ∂K. Let $S \subseteq \mathbb{R}^d$ be a finite set of sites. For each pair of sites $a, b \in S$, the *open halfsphere in the direction of $b - a$*, denoted by $H(b-a)$, is the set of points in ∂K whose exterior normal cone is contained in $(b-a)^\vee := \{\lambda \mid \lambda(b-a) > 0\}$. Intuitively, if we illuminate K with a light source at infinity in the direction $b - a$, then $H(b-a)$ is the the shadow of K, excluding its border.

For a fixed site $a \in S$, the *sector of a* is the set

$$H_S(a) \;=\; \bigcap_{b \in S \setminus \{a\}} H(b-a) \;.$$

We denote $\mathcal{H}_S := \{H_S(a) \mid a \in S\}$. Observe that $H(b-a)$ and, hence, $H_S(a)$, are open in ∂K.

Lemma 2.3.13. *Let $F_1, \dots, F_m$ be the facets of K and let $\lambda_{F_i}(x) \le 1$ be the valid linear inequality defining F_i. Then, for each $a \in S$,*

$$H_S(a) \;=\; \operatorname{relint}\left(\bigcup \{F_i \mid \; \lambda_{F_i}(a) < \lambda_{F_i}(b) \text{ for } b \in S \setminus a\}\right) \;.$$

In particular, $H_S(a) \cap H_S(b) = \emptyset$ for every $a, b \in S$ and, if S is in weak general position,

$$\bigcup_{a \in S} \overline{H_S(a)} = \partial K .$$

where $\overline{H_S(a)}$ denotes the topological closure of $H_S(a)$.

Proof. It is clear from the definition that $H_S(a)$ contains the relative interior of every F_i with $\lambda_{F_i}(a) < \lambda_{F_i}(b)$ for $b \in S \setminus a$. By convexity of the cones $\{\lambda \mid \lambda(b-a) > 0\}$ for each a, b, $H(b-a)$ (hence $H_S(a)$) also contains the relative interior of every lower dimensional face contained only in such facets. This proves the first formula. The second part follows from the first and the fact that in weak general position the minimum of each λ_{F_i} is attained at a single point of S. □

Remark 2.3.14. *Assuming weak general position, Lemma 2.3.13 allows us to think of $\mathcal{H}_S$ as a labeling of the facets of K by the elements of S or, equivalently, as a map $\mathcal{F}(K) \to S$. If K is centrally symmetric, then each pair of opposite facets F and $-F$ belong one to $H(b-a)$ and the other to $H(a-b)$. If K is not, we can still guarantee that $H(a-b)$ is never empty, and always disjoint from $H(b-a)$. As a consequence, $H(b-a)$ (and hence $H_S(a)$) cannot contain all the (relative interiors of) facets of K.*

For the case $K = \mathbb{B}^d$ of the tropical ball this partition of the facets translates into something more meaningful. Recall (see Proposition 2.3.11 and the paragraph before it) that facets of $\mathbb{B}^d$ can be represented as the arcs in the complete digraph on $d+1$ nodes. $\mathcal{H}_S$ colors these arcs (or facets) by the points of S. Then:

1. *The arrows in each color class form a partial order on the vertices of the complete digraph with $d+1$ vertices. This is, there is no monochromatic cycle.*

2. *For the case of two points in general position, the two colors are opposite acyclic tournaments. In particular, there is a bijection between the possible halfspheres $H(b-a)$ and the total orderings of $d+1$ elements.*

Theorem 2.3.15 ([29] for $d = 3$)**.** *Let $a, b \in \mathbb{R}^d$ be in weak general position. Then the central projection from a induces a homeomorphism between* $\mathrm{bisector}(a, b)$ *and $a + H(b-a)$. Hence,* $\mathrm{bisector}(a, b)$ *is homeomorphic to $\mathbb{R}^{d-1}$.*

Proof. Let us first show that $\mathrm{bisector}(a, b)$ is contained in $a + \mathrm{cone}(H(b-a))$. To seek a contradiction, let $c \in \mathrm{bisector}(a, b)$ such that $c - a \notin \mathrm{cone}(H(b-a))$. This implies that the smallest ball centered at a that contains c touches it at a facet F with functional λ_F such that $\lambda_F(b-a) \leq 0$. Now, c is equidistant to a and b, and a and b cannot be in the same facet of the ball centered at c (because they are in weak general position). Therefore, $\mathrm{dist}(c, \varepsilon a + (1-\varepsilon)b) < \mathrm{dist}(c, a)$, by convexity of the ball. This contradicts the fact that $\lambda_{F_i}(b-a) \leq 0$.

Hence, we have a well-defined map $\phi : \mathrm{bisector}(a, b) \to a + H(b-a)$ given by central projection. The map ϕ is continuous since it is the restriction of central projection. It is also proper (that is, the inverse image of a compact set is compact) by a following argument: Let C be a compact subset of $H(b-a)$. By continuity, $\phi^{-1}(C)$ is closed in $\mathrm{bisector}(a, b)$, hence in $\mathbb{R}^d$ since $\mathrm{bisector}(a, b)$ itself is closed (it is the zero set of the continuous function $d(x,a) - d(x,b)$). Thus, we only need to prove that $\phi^{-1}(C)$ is bounded. This follows from the fact that

$$\phi^{-1}(C) \subset (a + \mathrm{cone}(K)) \ \cap \ (b + \mathrm{cone}(\overline{H(a-b)})) \ ,$$

and that $\mathrm{cone}(C)$ and $\mathrm{cone}(\overline{H(a-b)})$ are two closed linear cones meeting only at the origin, since $H(a-b)$ and $H(b-a)$ are open and disjoint in ∂K.

Once we know ϕ is proper and continuous, we only need to check that it is bijective in order for it to be a homeomorphism. To show this, we construct its inverse. For each $v \in H(b-a)$ we consider the ray $r_v = \{a + \alpha v : \alpha \geq 0\}$. Along r_v, the distance to a is linear in α, the distance to b is convex in α and both functions are continuous. Observe also that

$$\mathrm{dist}(a+0v, a) = 0\,, \quad \mathrm{dist}(a+0v, b) > 0\,, \text{ and } \lim_{\alpha\to\infty} \mathrm{dist}(a+\alpha v, b) < \lim_{\alpha\to\infty} \mathrm{dist}(a+\alpha v, a) \ .$$

The last inequality comes from the fact that as the we move farther away from a along r_v, eventually $(a + \alpha v) - b$ will be in the same cone of $\mathcal{F}(K)$ as $(a + \alpha v) - a = \alpha v$ (by weak general position), and $\langle b - a, v\rangle > 0$ since $v \in H(b-a)$.

Hence, the function $\alpha \mapsto \mathrm{dist}(a + \alpha v, b) - \mathrm{dist}(a + \alpha v, a)$ is negative at zero, positive at infinity, continuous, and convex. Therefore, it has exactly one root, which means r_v intersects the bisector exactly once. We define $\psi(a+v)$ as this unique intersection point.

The maps ϕ and ψ are clearly inverses of one another. $\square$

Looking at the proof, the reader can check that central projection gives a proper and continuous map $\mathrm{bisector}(a,b) \to a + H(b-a)$ even without assuming weak general position. We only need weak general position to construct its inverse.

Corollary 2.3.16. *If S is in weak general position and there is an empty $H(a) \in \mathcal{H}_S$ then the bisector of S is empty.*

Proof. Assume that there is a point $c \in \mathrm{bisector}(S)$. For a site $a \in S$ let us show that $H_S(a) \neq \emptyset$. By definition, $c \in \mathrm{bisector}(a,b)$ for $b \in S \setminus \{a\}$. By Proposition 2.3.15, each $\mathrm{bisector}(a,b)$ can be mapped to $H(b-a)$ by central projection. Since c is in all of these bisectors, the central projection of c into the ball $a + K$ lies in $H(b-a)$ for all b, and hence in $H_S(a)$. $\square$

The converse of Proposition 2.3.16 is true for *three* points in arbitrary dimension (Theorem 2.3.18) but not for more, even in general position, as the following example shows:

Example 2.3.17 (Empty bisector, with non-empty sectors)**.** *Let $a = (1,-1,0,0)$, $b = (-1,1,0,0)$, $c = (0,0,2,-2)$ and $d = (0,0,-2,2)$. Then we have*

$$\mathrm{bisector}(a,b) = \{x \,|\, x_3 + 1 \leq x_1, x_2, \leq x_4 - 1\} \cup \{x \,|\, x_4 + 1 \leq x_1, x_2, \leq x_3 - 1\} \cup \{x \,|\, x_1 = x_2\} \ .$$

By symmetry, we also have

$$\mathrm{bisector}(c,d) = \{x \,|\, x_1 + 2 \leq x_3, x_4, \leq x_2 - 2\} \cup \{x \,|\, x_2 + 2 \leq x_3, x_4, \leq x_1 - 2\} \cup \{x \,|\, x_3 = x_4\} \ .$$

Since $\mathrm{bisector}(a,b,c,d)$ *lies in the intersection of the two, we have*

$$\mathrm{bisector}(a,b,c,d) \ \subseteq \ \{x \,|\, x_1 = x_2, x_3 = x_4\} \ .$$

So for $x \in \text{bisector}(a,b,c,d)$*, we may assume* $x_3 = 0$*, which entails:*

$$\text{dist}(a,x) \;=\; \max\{x_1+1,0\} - \min\{x_1-1,0\} \;=\; \max\{|x_1|+1,2\} \le |x_1|+2 \;\; ,$$

with equality only when $x_1 = 0$ *and*

$$\text{dist}(c,x) \;=\; \max\{x_1,2\} - \min\{x_1,-2\} \;=\; \max\{|x_1|+2,4\} \ge |x_1|+2 \;\; ,$$

with equality only if $|x_1| \ge 2$*. This shows that* $\text{bisector}(a,b,c,d) = \emptyset$*.*

However, $\mathcal{H}_S$ *has no empty sector since the sectors of* a,b,c,d *contain the facets with outer normals* $(1,-1,0,0)$*,* $(-1,1,0,0)$*,* $(0,0,1,-1)$ *and* $(0,0,-1,1)$*. Both this property and the emptiness of* $\text{bisector}(a,b,c,d)$ *are preserved under a small perturbation; cf. Remark 2.3.2.*

2.3.3 Bisectors of three points

The goal of this section is to prove our first main result.

Theorem 2.3.18. *Let* $S = \{a_1,a_2,a_3\}$ *be a set of three distinct points in* $\mathbb{R}^d$ *which lie in weak general position with respect to a convex body* K*. If* $H_S(a_i) \neq \emptyset$ *for* $i = 1,2,3$ *then* $\text{bisector}(a_1,a_2,a_3)$ *is homeomorphic to a non-empty open subset of* $\mathbb{R}^{d-2}$*.*

Corollary 2.3.19. *For any three points in weak general position* $\text{bisector}(a_1,a_2,a_3)$ *is either empty or pure of dimension* $d-2$*. If* $d = 2$ *then* $\text{bisector}(a_1,a_2,a_3)$ *is either empty or a single point.*

We show first the two-dimensional case of Theorem 2.3.18. It was used in [29]:

Lemma 2.3.20. *Let* $a_1,a_2,a_3 \in \mathbb{R}^2$ *be in weak general position with respect to a convex body* K*. If* $H_S(a_i) \neq \emptyset$ *for the three of them, then* $\text{bisector}(a_1,a_2,a_3)$ *is a point.*

Proof. Suppose first that $\text{bisector}(a_1,a_2,a_3)$ is empty; i.e., the three two-point bisectors do not meet. Then one of them, say $\text{bisector}(a_1,a_3)$, does not appear at all in Vor_S. We will show that this implies $H_S(a_2) = \emptyset$.

To simplify the exposition, we assume that the line a_1a_3 is horizontal. Let u and v be the points where the ball K has a horizontal tangent. Observe that u and v are unique, by weak general position. By Theorem 2.3.15, $\text{bisector}(a_1,a_3)$ is a connected curve having as asymptotic directions those of u and v.

We are assuming that $\text{bisector}(a_1,a_3) \subset \text{Vor}_S(a_2)$. That is, every ball $x + \alpha K$ with a_1 and a_3 in the boundary (hence with center $x \in \text{bisector}(a_1,a_3)$) has a_2 in the interior. As we move x towards infinity in the direction of u, these balls converge towards the translation of the cone $\text{cone}(K-v)$ to have a and b in the boundary. As we move towards v, the balls converge towards a translation of $\text{cone}(K-u)$. Thus, a_2 lies in the interior of these two cones, which implies $H_S(a_2)$ to be empty.

We conclude that, if $H_S(a_i) \neq \emptyset$ for $i = 1,2,3$ then $\text{bisector}(S) \neq \emptyset$. It remains to show that $\text{bisector}(S)$ is a unique point. This is equivalent to saying that the 1-parameter family of balls having a_1 and a_3 in the boundary contains a unique element with a_2 in the boundary. To show this, suppose without loss of generality that a_2 is above the line a_1a_3. Then, for every $x \in \text{bisector}(S)$ we have that when we move x up along $\text{bisector}(a_1,a_3)$,

a_2 enters the interior of the ball centered at x and with a_1, a_3 in the boundary. Since this happens for *every* $x \in \text{bisector}(S)$, it can only happen once as we move x along $\text{bisector}(a_1, a_3)$. $\square$

For the rest of the proof let $S = \{a_1, a_2, a_3\}$ be three points in general position with respect to a polytopal convex body K. We turn the general problem into a two-dimensional one via the following construction: Let $\pi_S : \mathbb{R}^d \to \mathbb{R}^{d-2}$ be the affine projection that quotients out the 2-plane Π containing S. We first show some properties of the map π:

Lemma 2.3.21. *With the above notation, let $x \in \text{int}(\pi(K)) \subset \mathbb{R}^{d-2}$. Let $\Pi_x := \pi^{-1}(x)$ (a 2-plane parallel to Π) and let $K_x = K \cap \Pi_x$. Then K_x is a convex polygon.*

Proof. If K_x is not a convex, full dimensional polygon, it must be either an edge or a vertex of K. This is because K_x is not empty (by definition of K_x), and it is a polytope because K_x is a section of K by a 2-dimensional affine subspace. In any case, it is a face, and parallel to Π, so x has to be in the frontier of $\pi(K)$, which contradicts the assumption. For this same reason, if $x \in \text{int}(\pi(K))$, then K_x cannot be a facet of K. $\square$

Lemma 2.3.22. *Let K, $x \in \text{int}(\pi(K))$ as before. Let $H_S^{(x)}(a_i)$ denote the sector of a_i computed with respect to K_x (which is a polygon according to the previous lemma). Then:*

$$H_S^{(x)}(a_i) \;=\; H_S(a_i) \cap K_x \;\; .$$

Proof. Note that no facet in K_x will be parallel to any $a_i - a_j$ because if it were, the corresponding facet in K would be parallel too. Then, $H_S^{(x)}$ is well defined.

Let F' be a facet of K_x, and let F be the corresponding facet in K. Then, $n(F') \in \mathbb{R}^2$, the normal vector to F', is the projection into Π of $n(F)$, and,

$$\begin{aligned} F' \in H_S^{(x)}(a_i) \quad &\iff \quad \langle n(F'), a_i \rangle > \langle n(F'), a_j \rangle \text{ for } j = 1,2,3 \\ &\iff \quad \langle n(F), a_i \rangle > \langle n(F), a_j \rangle \text{ for } j = 1,2,3 \quad \iff \quad F \in H_S(a_i). \end{aligned}$$

$\square$

Lemma 2.3.23. *Let $S = \{a_1, a_2, a_3\}$ as before. If $H_S(a_i) \neq \emptyset$ for all i, then $\bigcap_{a_i \in S} \pi(H_S(a_i))$ is open and not empty.*

Proof. First, observe that an $x \in \partial K$ with $\pi(x) \in \partial(\pi(K))$ cannot be in any of the $H_S(a_i)$: indeed, $x \in \partial K$ implies that there is a normal vector of K at x orthogonal to Π, hence orthogonal to $a_i - a_j$ for every a_i, a_j. As a consequence,

$$\bigcap_{a_i \in S} \pi(H_S(a_i)) \;\subset\; \text{int}\, \pi(K) \;\; ,$$

which implies it is open.

For any point $x \in \text{int}(\pi(K))$, the preimage $\pi^{-1}(x)$ is a polygon, a slice of K. This slice has to intersect at least two of the classes of $\mathcal{H}$, because $\mathcal{H}$ is a partition (so at the slice intersects at least one class), but no class can contain a set of facets whose vectors are positively dependent (because each class is an intersection of half-spheres). Then, any point $x \in \text{int}(K)$ lies in at least two sets $\pi(H_S(a_i))$.

Thus, the three open sets $\pi(H_S(a_i))$ cover each point of $\mathrm{int}(\pi(K))$ at least twice. At least two of these sets must intersect, say $\pi(H_S(a_1))$ and $\pi(H_S(a_2))$. Suppose that $(\pi(H_S(a_1)) \cap \pi(H_S(a_2))) \cap \pi(H_S(a_3)) \neq \emptyset$. Then, $\mathrm{int}(\pi(K))$ would be disconnected, because it is covered by two disjoint open sets. Since this is not possible, there must be a point in the common intersection of the three $H_S(a_i)$. $\square$

Proof of Theorem 2.3.18. Consider the map

$$\phi : \mathrm{bisector}(a_1, a_2, a_3) \longrightarrow \bigcap_{i=1,2,3} \pi(H_S(a_i))$$

defined as follows: Let $p \in \mathrm{bisector}(S)$, and let $v_i \in H_S(a_i)$ be the central projection from p to $a_i + \partial K$, for each $i = 1, 2, 3$. Note that each v_i lies in the corresponding $H_S(a_i)$, by Theorem 2.3.15. Further, the three points v_1, v_2 and v_3 lie in a plane parallel to Π. In particular, $\pi(v_1) = \pi(v_2) = \pi(v_3)$ lies in $\bigcap_{i=1,2,3} \pi(H_S(a_i))$ and we define

$$\phi(p) \ := \ \pi(v_i) \ .$$

To show that ϕ is a homeomorphism, let us construct its inverse ψ. Let $\gamma : \pi(\mathrm{int}(K)) \to \mathrm{int}(K)$ be a continuous section of π in K. For example, but not necessarily, for each 2-plane Π' parallel to Π and intersecting K let $\gamma(\pi(\Pi'))$ be the centroid of $\Pi' \cap K$.

Now, let $x \in \bigcap_{i=1,2,3} \pi(H_S(a_i))$. Let $\Pi_x = \pi^{-1}(x)$ and let $w_i = \gamma(x) + a_i$, for each $i = 1, 2, 3$. In the 2-plane Π_x we have a set $S_x = \{w_1, w_2, w_3\}$ and a unit ball $K \cap \Pi_x$. Lemma 2.3.21 gives that $H_{S_x}(w_i) = H_S(a_i) \cap \Pi_x$. By choice of x we have $\bigcap_i H_{S_x}(w_i) \neq \emptyset$ and Lemma 2.3.20 guarantees that the bisector of S_x is a unique point $r \in \Pi_x$; cf. Figure 2.3.

Let v_i be the central projection of r to $w_i + \partial K_x$. Observe that $|w_i r|/|w_i v_i|$ is independent of i since it equals $\mathrm{dist}_{K_x}(w_i, r)$, where dist_{K_x} denotes the distance induced by K_x in $\mathbb{R}^2$, and r is in $\mathrm{bisector}_{K_X}(w_1, w_2, w_3)$. Since $w_i - a_i = \gamma(x)$ is also independent of i, the three rays $a_i v_i$ meet at the point

$$p \ = \ r + \mathrm{dist}_{K_x}(w_i, r)\gamma(x) \ ,$$

and $\mathrm{dist}_K(a_i, p) = \mathrm{dist}_{K_X}(w_i, r)$. Thus, $p \in \mathrm{bisector}(a_1, a_2, a_3)$ and we define p to be $\psi(x)$.

This gives us a well-defined map

$$\psi : \bigcap_{i=1,2,3} \pi(H_S(a_i)) \longrightarrow \mathrm{bisector}(a_1, a_2, a_3) \ ,$$

and by construction ψ is the inverse of ϕ both ways. The map γ is continuous, and the bisector of three points in the plane depends continuously on a continuous deformation of the unit ball. It follows that ψ is continuous. Thus, ϕ and ψ are homeomorphisms between $\mathrm{bisector}(a_1, a_2, a_3)$ and $\bigcap_{i=1,2,3} \pi(H_S(a_i))$. Since the latter is not empty and open by Lemma 2.3.23, the former is homeomorphic to a non-empty open subset of $\mathbb{R}^{d-2}$. $\square$

Our proof of Theorem 2.3.18 closely follows the proof of the 3-dimensional case in [30]. There, it is additionally shown that the number of connected components of

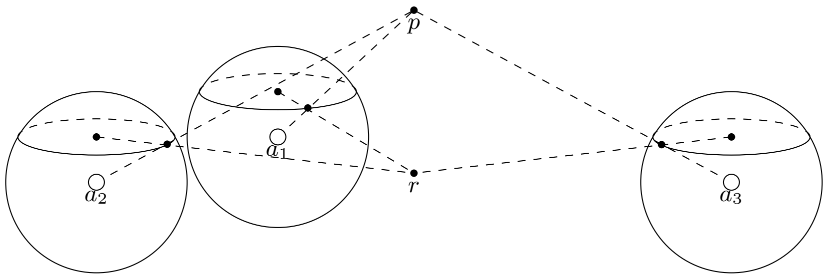

Figure 2.3: Construction of ψ in the proof of Theorem 2.3.18

bisector(a_1, a_2, a_3) equals the total number of connected components of the three sectors $H_S(a_1)$ minus two. We can extend this to higher dimension and to higher reduced Betti numbers (over any field):

Theorem 2.3.24. *Let $a_1, a_2, a_3 \in \mathbb{R}^d$ be three points in weak general position with respect to a convex body K and assume that $H_S(a_i) \neq \emptyset$ for all three. Then, for $j \in \{0, \dots, d-3\}$, we have*

$$\tilde{\beta}_j(\mathrm{bisector}_K(a_1, a_2, a_3)) \;=\; \sum_{i=1}^{3} \tilde{\beta}_j(H_S(a_i)) \ .$$

Proof. Consider the same projection $\pi : \mathbb{R}^d \to \mathbb{R}^{d-1}$ as before. Observe that, since $\Pi_x \cap H_S(a_i)$ is empty or contractible for every plane Π_x parallel to Π, we have

$$H_S(a_i) \;\simeq\; \pi(H_S(a_i)) \text{ for } i = 1, 2, 3 \ .$$

We now apply Alexander duality in the one point compactification $\mathcal{S}$ of $\mathrm{int}(\pi(K))$, which is a sphere of dimension $d-2$. Alexander duality says that if U is an open and locally contractible subset of a sphere $\mathcal{S}$ then

$$\tilde{\beta}_j(U) \;=\; \tilde{\beta}_{d-3-j}(\mathcal{S} \setminus U) \ .$$

In particular, if we let $C_i = \mathcal{S} \setminus \pi(H_S(a_i))$ we have

$$\sum_{i=1}^{3} \tilde{\beta}_j(\pi(H_S(a_i))) \;=\; \sum_{i=1}^{3} \tilde{\beta}_{d-3-j}(C_i) \ ,$$

and

$$\tilde{\beta}_j\left(\bigcap_{i=1}^{3} \pi(H_S(a_i))\right) \;=\; \tilde{\beta}_{d-3-j}\left(\bigcup_{i=1}^{3} C_i\right) \ .$$

Yet C_1, C_2 and C_3 are pairwise disjoint except for the "point at infinity" of $\mathcal{S}$, because each point of $\mathrm{int}(\pi(K))$ lies in at least two of the sets $\pi(H_S(a_i))$. Thus, $\bigcup_{i=1}^{3} C_i$ is the topological wedge (or one-point sum) of C_1, C_2 and C_3, which makes the right-hand sides of the two last equations coincide. $\square$

Remark 2.3.25. *One may ask how complicated the Betti numbers $\tilde{\beta}_j(H_S(a_i))$ in Theorem 2.3.24 can be. Equivalently, how complicated the topology of three point bisectors can* *be. Such a bisector is $(d-2)$-dimensional, so the relevant Betti numbers are $\tilde\beta_0, \ldots, \tilde\beta_{d-2}$. The last one, $\tilde\beta_{d-2}$, must vanish as $\tilde\beta_j(H_S(a_i)) = \tilde\beta_j(\pi(H_S(a_i)))$, and the latter is an open subset of $\mathbb{R}^{d-2}$. But $\tilde\beta_{d-3}$ can be non-zero, as the following example of a disconnected bisector of three points for the tropical ball in dimension three shows.*

Example 2.3.26. *Consider the points $a = (0,0,4,4)$, $b = (-3,0,2,0)$ and $c = (0,-3,0,2)$ in weak general position in $\mathbb{R}^4/\mathbb{R}\mathbb{1}$. We are going to describe the three sectors. For each of them we list the facets whose relative interior is in the corresponding sector. The facet F_{ij} is the one at which coordinate i is minimized and j is maximized.*

$$H_a = (F_{14}, F_{23})\,, \quad H_b = (F_{12}, F_{13}, F_{32}, F_{42}, F_{43})\,, \quad H_c = (F_{21}, F_{24}, F_{31}, F_{34}, F_{41})\ .$$

The sector H_a is not connected and, hence, bisector a, b, c *is not connected.*

2.4 Classification of tropical bisectors of two points

2.4.1 Tropical bisectors and tropical hypersurfaces

In the classical case the bisector of two points is a degenerate quadric, namely the affine hyperplane perpendicular to the connecting line segment and which runs through the midpoint. The tropical analog is more interesting.

Proposition 2.4.1. *Let $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ be in weak general position. Then the homogeneous* max*-tropical Laurent polynomial*

$$\phi(a,b) \;=\; \max\left(\max_{i,j\in[d+1]}(x_i - a_i - x_j + a_j), \max_{k,\ell\in[d+1]}(x_k - b_k - x_\ell + b_\ell)\right) \tag{2.6}$$

vanishes on bisector(a,b)*. That is, the set* bisector(a,b) *is contained in a* max*-tropical hypersurface of degree $d+1$.*

Proof. Recall that a max-tropical (Laurent) polynomial vanishes if the maximum is attained at least twice; cf. [52, §3.1]. First, we check that there are no duplicates among the terms in the representation (2.6) of $\phi(a,b)$. Assume the contrary, i.e., $x_i - a_i - x_j + a_j = x_i - b_i - x_j + b_j$ for some $i, j \in [d+1]$. Then $a_j - a_i = b_j - b_i$, which forces $\langle e_j - e_i, b - a\rangle = (b_j - a_j) - (b_i - a_i) = 0$. Thus $b - a$ is parallel to the facet of $\mathbb{B}^d$ with normal vector $e_j - e_i$. This was explicitly excluded in our assumption, and we arrive at the desired contradiction. We infer that the $2d(d+1)$ terms are pairwise distinct.

Let $x \in \operatorname{bisector}(a,b)$. This means that $\operatorname{dist}(a,x) = \operatorname{dist}(b,x)$, and thus $\max_{i,j\in[d+1]}(x_i - a_i - x_j + a_j) = \max_{k,\ell\in[d+1]}(x_k - b_k - x_\ell + b_\ell)$. It follows that $\phi(a,b)$ vanishes at x.

The degree of the bisector tropical hypersurface can be read off any Laurent monomial like $x_i - x_j$ by adding $x_1 + x_2 + \cdots + x_{d+1}$, which yields the true monomial $2x_i + \sum_{k\in[d+1]-\{i,j\}} x_k$. The latter has degree $2 + (d+1-2) = d+1$. □

Proposition 2.4.1 yields a trivial algorithm to compute tropical bisectors in weak general position: enumerate the maximal cells of the tropical hypersurface defined by (2.6) and select those maximal cells that attain maxima in one monomial of type $x_i - a_i - x_j + a_j$ and one monomial of type $x_k - b_k - x_\ell + b_\ell$. This algorithm needs to go through the $\Theta(d^4)$ choices of one monomial from the left and one from the right. This is worst case optimal, as we will prove in Corollary 2.4.12 that tropical bisectors can have $\Omega(d^4)$ maximal cells.

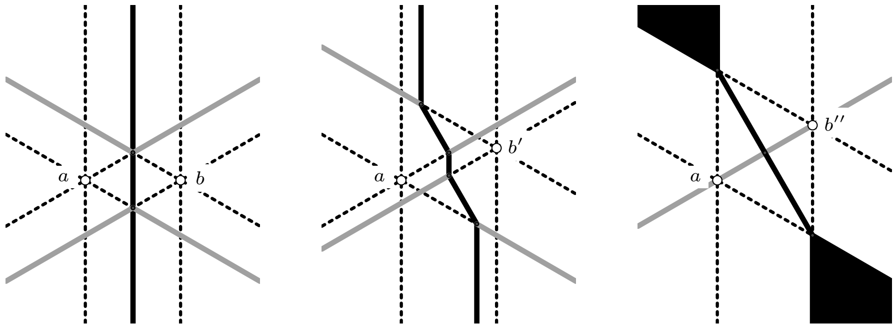

Figure 2.4: Tropical bisectors for $b-a=(-1,1,0)$, $(-1,1,\frac{1}{2})$ and $(-1,1,1)$, respectively. The first two are in weak general position but the last one is not. Only the middle one is in general position. The bisectors are shown in black and the rest of the tropical hypersurface containing it in gray; cf. Prop. 2.4.1.

Example 2.4.2. *The labeling of the faces of a tropical bisector does not need to be unique if a and b are not in weak general position. For instance, if $b-a=(-1,1,0)$ then*

$$\begin{aligned}\operatorname{bisector}_{(-+*),(+-*)}(a,b) &= \operatorname{bisector}_{(-++),(+-+)}(a,b) = \operatorname{bisector}_{(-++),(+--)}(a,b)\\ &= \operatorname{bisector}_{(-+-),(+-+)}(a,b) = \operatorname{bisector}_{(-+-),(+--)}(a,b)\end{aligned}$$

is the only face; see Figure 2.4 (left).

2.4.2 The bisection fan

Normal equivalence of tropical bisectors is preserved by translation and scaling. In particular the equivalence class of bisector(a,b) is uniquely determined by the direction of the vector $b-a$. The *bisection fan* $\mathcal{F}^d_{\mathrm{bis}}$ is the complete polyhedral fan in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ whose relatively open cones are defined by "x and y lie in the same cone if and only if bisector$(0,x)$ and bisector$(0,y)$ are normally equivalent". Put differently, two points $a,b\in\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ are in *general position* if, and only if, the difference $b-a$ lies in a maximal cone of $\mathcal{F}^d_{\mathrm{bis}}$. In the rest of this section we show that $\mathcal{F}^d_{\mathrm{bis}}$ is indeed a polyhedral fan and give an explicit description of it.

Recall that an *ordered partition* or *total preorder* on a finite set S is a partition of S into non-empty parts together with a total order on the parts. If the parts are denoted $S_1,\dots,S_k$ (in this order), we can write $x\le y$ meaning "$x\in S_i$ and $y\in S_j$ for some $i\le j$". In particular, for all $x,y\in S$ we have $x\le y\le x$ if and only if x and y lie in the same part.

Any real vector $v=(v_1,\dots,v_{d+1})\in\mathbb{R}^{d+1}$ induces an ordered partition $S(v)$ of $[d+1]$ by putting together the coordinates that have the same value and ordering the groups according to their values. For example, the vector $v=(3,1,6,4,6,3,1)$ of length seven induces the partition $(\{2,7\},\{1,6\},\{4\},\{3,5\})$ of the set $\{1,2,\dots,7\}$ into four parts.

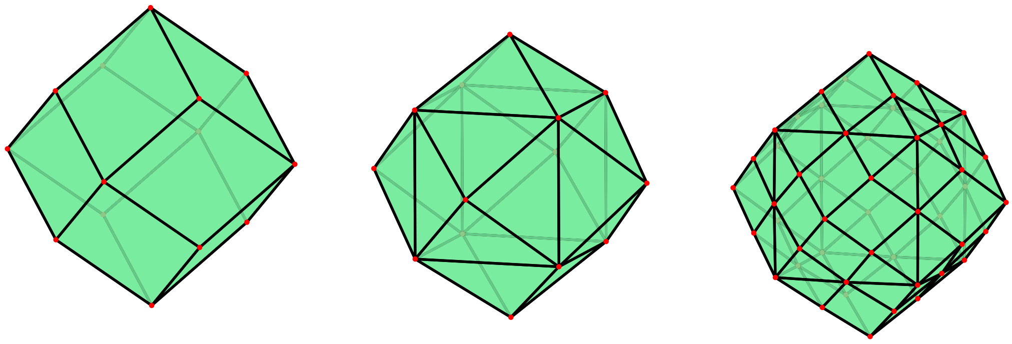

Figure 2.5: The fans $\mathcal{F}(\mathbb{B}^3)$, $\mathcal{F}(A_3)$, and the bisection fan $\mathcal{F}^3_{\mathrm{bop}}$.

Note that the ordered partition $S(v)$ is constant on the class $v + \mathbb{R}\mathbb{1}$. Hence these ordered partitions are defined for points in the projective tropical torus $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$.

For $v, w \in \mathbb{R}^{d+1}$ with $S(v) = S(w)$ we have $S(v + w) = S(v)$ and, moreover, $S(\alpha v) = S(v)$ for any positive real α. That is to say, the stratification of $\mathbb{R}^{d+1}$ by ordered partitions forms a complete polyhedral fan. In what follows we seek to refine that fan by recording which part, or which gap between parts, contains the *midvalue*

$$\mu(v) \ := \ \frac{1}{2}\big(\max_{i\in[d+1]} v_i + \min_{i\in[d+1]} v_i\big) \ .$$

The *bisected ordered partition* of $[d+1]$ induced by $v \in \mathbb{R}^{d+1}$ is the ordered partition $S(v)$ as defined above, together with the information of which (values of) parts are smaller, equal or greater than the midvalue $\mu(v)$; see also [21]. Equivalently, this is the ordered partition associated with the *extended vector* $(v, \mu(v)) \in \mathbb{R}^{d+2}$. We denote by $\mathcal{F}^d_{\mathrm{bop}}$ the fan of bisected ordered partitions of dimension d.

Remark 2.4.3. *The "finest" or "most generic" ordered partitions are the permutations, in which each part is a singleton. Hence, the fan of ordered partitions equals the normal fan of the* permutahedron. *This, in turn, coincides with the fan of regions in the* braid arrangement *or Coxeter arrangement of type A_d, which consists of the hyperplanes $\{x \mid x_i = x_j\}$ for $1 \le i < j \le d+1$. We denote this fan $\mathcal{F}(A_d)$. It is intermediate between the central fan of the tropical ball (which is coarser) and the fan of bisected ordered partitions (which is finer):*

$$\mathcal{F}(\mathbb{B}^d) \ \le \ \mathcal{F}(A_d) \ \le \ \mathcal{F}^d_{bop} \ ;$$

see Figure 2.5 for a visualization of the case $d = 3$. Note that $\mathcal{F}(A_d)$ is also the fan of weak general position: a and b are in weak general position if and only if $b - a$ lies in a full-dimensional cell.

Example 2.4.4. *In $d = 2$, the fans $\mathcal{F}(A_d)$ and $\mathcal{F}(\mathbb{B}^d)$ coincide, they form the face fan of the regular hexagon $\mathbb{B}^2$. The bisection fan $\mathcal{F}^d_{bop}$ is the barycentric subdivision of it; cf. Figure 2.6. Excluding permutations of the coordinates, and sign inversion, we infer that there are three types of tropical bisectors in the plane, and these are shown in Figure 2.4. The type to the left is in weak general position but not in general position, the type to the left is in general position, and the type to the right is not even in weak general position.*

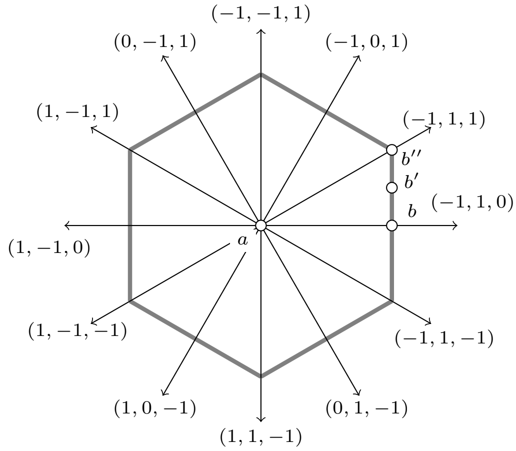

Figure 2.6: The bisection fan, $\mathcal{F}^d_{\text{bis}}$, for $d = 2$. The three vectors $b - a$ for Figure 2.4 have been marked.

Recall that the max-tropical *line segment* between two points, a and b, is the set

$$[a, b] \ := \ \{\max(\alpha\mathbb{1} + a, \beta\mathbb{1} + b) \mid \alpha, \beta \in \mathbb{R}\} \ \subset \ \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \ .$$

It is worth noting that the combinatorial types of tropical segments are classified by the braid fan $\mathcal{F}(A_d)$; cf. [52, Prop. 5.11]. In this sense the fan of bisected ordered partitions $\mathcal{F}^d_{\text{bop}}$ classifies combinatorial types of "bisected tropical segments":

Proposition 2.4.5. *The bisected ordered partition of $b-a$ contains the same information as the combinatorial type of the tropical line segment $[a, b]$ together with the information of which part contains the midpoint.*

Proof. Suppose for simplicity that $a, b \in \mathbb{R}^{d+1}$ satisfy

$$b_1 - a_1 \ \leq \ b_2 - a_2 \ \leq \ \ldots \ \leq \ b_{d+1} - a_{d+1} \ . \tag{2.7}$$

Then $[a, b]$ is the union of at most d ordinary line segments, one for each subset of coordinates between a strict inequality in (2.7). That is, the combinatorics of the tropical segment is the same as the ordered partition of $b - a$. The midvalue $\mu(b - a)$ selects one of the ordinary segments. □

The goal of the rest of this section is to prove the following:

Theorem 2.4.6. $\mathcal{F}^d_{bis} = \mathcal{F}^d_{bop}$. *That is, given $a, b, a', b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ we have* bisector(a, b) *is normally equivalent to* bisector(a', b') *if, and only if, $b - a$ and $b' - a'$ induce the same bisected ordered partition of* $[d + 1]$.

2.4.3 Proof of Theorem 2.4.6

In (2.3) we defined the type (F_-, F_*, F_+) of a face F of $\mathbb{B}^d$ or the corresponding face cone in $\mathcal{F}(\mathbb{B}^d)$. For a pair of faces, F and G, this gives rise to the following *labeling partition* of $[d+1]$:

$$\begin{aligned} L_0 &:= (F_- \cap G_-) \cup (F_+ \cap G_+) \ , \\ L_+ &:= \big(F_+ \cap G_*\big) \cup \big(F_* \cap G_-\big) \ , \\ L_- &:= \big(F_- \cap G_*\big) \cup \big(F_* \cap G_+\big) \ , \\ L_{+1} &:= F_+ \cap G_- \ , \\ L_{-1} &:= F_- \cap G_+ \ , \\ L_* &:= F_* \cap G_* \ . \end{aligned} \tag{2.8}$$

As a first step in the proof of Theorem 2.4.6, the following lemma characterizes when is $\mathrm{bisector}_{(F,G)}(a,b)$ non-empty. Recall that this is the case if and only if there is a tropical ball touching a and b at faces F and G, respectively.

Lemma 2.4.7. *Let F and G be a fixed pair of faces of $\mathbb{B}^d$ with the labeling partition defined as in* (2.8)*. Further let $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. Then the set* $\mathrm{bisector}_{(F,G)}(a,b)$ *is not empty if, and only if, there exist $\gamma \in \mathbb{R}$ and $\delta \in [0,\infty)$ such that the following conditions are satisfied:*

$$\begin{cases} (b-a)_i = \gamma & \textit{if } i \in L_0 \ , \\ (b-a)_i \in [\gamma, \gamma+\delta] & \textit{if } i \in L_+ \ , \\ (b-a)_i \in [\gamma-\delta, \gamma] & \textit{if } i \in L_- \ , \\ (b-a)_i = \gamma - \delta & \textit{if } i \in L_{-1} \ , \\ (b-a)_i = \gamma + \delta & \textit{if } i \in L_{+1} \ , \\ (b-a)_i \in [\gamma-\delta, \gamma+\delta] & \textit{if } i \in L_* \ . \end{cases} \tag{2.9}$$

Proof. Let us assume that the face of the tropical bisector $\mathrm{bisector}(a,b)$ defined by (F,G) is non-empty. Then there is a point x such that $\mathrm{dist}(a,x) = \mathrm{dist}(x,b) = \delta$ and $a - x \in F$ as well as $b - x \in G$. We set $\gamma_a = \min_{i\in[d+1]}(a_i - x_i)$ and $\gamma_b = \min_{i\in[d+1]}(b_i - x_i)$. The possible values for the coordinates of $a - x$ and $b - x$ are

$$\begin{cases} (a-x)_i = \gamma_a + \delta & \text{if } i \in F_- \ , \\ (a-x)_i = \gamma_a & \text{if } i \in F_+ \ , \\ (a-x)_i \in [\gamma_a, \gamma_a+\delta] & \text{if } i \in [d+1] \setminus (F_- \cup F_+) \end{cases} \tag{2.10}$$

and

$$\begin{cases} (b-x)_i = \gamma_b + \delta & \text{if } i \in G_- \ , \\ (b-x)_i = \gamma_b & \text{if } i \in G_+ \ , \\ (b-x)_i \in [\gamma_b, \gamma_b+\delta] & \text{if } i \in [d+1] \setminus (G_- \cup G_+) \ , \end{cases} \tag{2.11}$$

for some $\delta \geq 0$. Setting $\gamma = \gamma_b - \gamma_a$ the above translates into (2.9).

For the converse, note that going from (2.10) and (2.11) to (2.9) is the Fourier-Motzkin elimination of the variables x_i. Therefore, any γ and $\delta \geq 0$ which are feasible for (2.9) can be lifted to a solution of (2.10) and (2.11). That is to say, we can set $\gamma_a = 0$ and $\gamma_b = \gamma$, and the conditions in (2.10) and (2.11) yield a point $x \in \mathrm{bisector}_{(F,G)}(a,b)$. □

Proposition 2.4.8. *Let F and G be a fixed pair of faces of $\mathbb{B}^d$. Then the set*

$$C \ := \ \big\{ b - a \ \big| \ a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \ \textit{with} \ \mathrm{bisector}_{(F,G)}(a,b) \neq \emptyset \big\} \tag{2.12}$$

is both a polyhedral cone and a tropical cone, although perhaps not a tropical polyhedral cone.

Proof. Let $a, a', b, b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ such that $\operatorname{bisector}_{(F,G)}(a,b)$ and $\operatorname{bisector}_{(F,G)}(a',b')$ both are nonempty. Since $\operatorname{bisector}_{(F,G)}(a,b) \neq \emptyset$, by Lemma 2.4.7, there are scalars γ and δ satisfying the conditions (2.9). Likewise there are certificates γ' and δ' for $\operatorname{bisector}_{(F,G)}(a',b') \neq \emptyset$. By linearity of the conditions (2.9) it follows that $\gamma+\gamma'$ and $\delta+\delta'$ certify that $\operatorname{bisector}_{(F,G)}(a+a', b+b') \neq \emptyset$: for instance, we have $(b+b'-a-a')_i = \gamma+\gamma'$ for $i \in L_0 = (F_- \cap G_-) \cup (F_+ \cap G_+)$. Since clearly $\alpha c \in C$ for all $c \in C$ and $\alpha \geq 0$ we conclude that C is an ordinary cone. This cone is polyhedral because it is defined in terms of the finitely many linear conditions (2.9).

A similar argument shows that C is also closed with respect to taking arbitrary $(\max, +)$-linear combinations: for instance, with the above notation we have $\max((b-a)_i, (b'-a')_i) = \max(\gamma, \gamma')$ for $i \in L_0$. This shows that C is a tropical cone. ☐

Corollary 2.4.9. *The bisection fan of tropical bisectors is a classical polyhedral fan, and a tropical (perhaps not tropical polyhedral) fan.*

Proof. We know that the feasibility region of a face (F, G) is a tropical and classical polyhedral cone. Finite intersections of these cones are again tropical cones, classical cones, and polyhedral cones. Therefore, the feasibility region of a normal equivalence class, which is the intersection of the cones of its non-empty faces, is again a tropical and classical polyhedral cone. Hence, the whole fan has this structure. ☐

The following shows one direction of Theorem 2.4.6, namely, $\mathcal{F}^d_{\text{bis}}$ is coarser than $\mathcal{F}^d_{\text{bop}}$.

Lemma 2.4.10. *Let $F, G \in \mathcal{F}(\mathbb{B}^d)$. Then, whether* $\operatorname{bisector}_{(F,G)}(a,b)$ *is empty or not, for each $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ depends only on the bisected ordered partition of $b-a$.*

Proof. Consider the partition of $[d+1]$ into six sets $L_0, L_+, L_-, L_{+1}, L_{-1}, L_*$ defined in (2.8). We want to show that feasibility of the system (2.9) for a given a and b depends only on the bisected ordered partition of $b-a$. Without loss of generality we assume $a = 0$ and $b_1 \leq b_2 \leq \cdots \leq b_{d+1}$ as in (2.7).

Let $\mu(b) := \frac{1}{2}(b_{d+1}+b_1) = \frac{1}{2}\operatorname{dist}(0,b) + b_1$ be the midvalue of $b-0$ and $m = \max(0+\mu(b)\mathbb{1}, b)$ the midpoint of the segment $[0,b]$. In particular, $\operatorname{dist}(0,m) = \operatorname{dist}(m,b) = \frac{1}{2}\operatorname{dist}(0,b)$.

We distinguish three cases, depending on whether both, none, or exactly one of L_{+1} and L_{-1} are empty.

Claim I: *Suppose that $L_{+1} \cup L_{-1} = \emptyset$. Then,* (2.9) *is feasible if, and only if,*

$$\textit{there are } k \leq \ell \textit{ with } L_- \subseteq \{1,\ldots,k\},\ \ L_0 \subseteq \{k+1,\ldots,\ell\},\ \ L_+ \subseteq \{\ell+1,\ldots,d+1\}\ . \tag{I.1}$$

Indeed, in this case feasibility of (2.9) is equivalent to feasibility of

$$\begin{cases} \gamma \geq b_i & \text{for } i \in L_-\ , \\ \gamma = b_i & \text{for } i \in L_0\ , \\ \gamma \leq b_i & \text{for } i \in L_+\ , \end{cases} \tag{2.13}$$

which implies the ordered partition to satisfy (I.1). Conversely, if the ordered partition satisfies (I.1), then let γ be chosen to satisfy (2.13) and let $\delta \geq \min(\gamma - b_1, b_{d+1} - \gamma)$. This yields a feasible solution to (2.9). In other words, in this case we can tell if C is empty or not by just looking at the the ordered partition of b; the relative position of the midpoint is irrelevant.

Claim II: *Suppose that $L_{+1} \neq \emptyset = L_{-1}$. Then,* (2.9) *is feasible if, and only if, in addition to* (I.1)*, we have*

$$\begin{aligned} \{b_i \mid i \in L_{+1}\} &= \{b_{d+1}\} \ , && \text{(II.1)} \\ b_i \leq \mu(b) \ \textit{for} \ i &\in L_0 \cup L_- \ , && \text{(II.2)} \\ |\{b_i \mid i \in L_0\}| &\leq 1 \ . && \text{(II.3)} \end{aligned}$$

Indeed, if (I.1), (II.1), (II.2), and (II.3) hold, then take $\gamma = \max\{b_i \mid i \in L_0 \cup L_- \cup \{1\}\}$ and $\delta = b_{d+1} - \gamma$.

Conversely, if (γ, δ) is feasible for (2.9) then $\gamma + \delta = b_i = b_{d+1}$ for all $i \in L_{+1}$. In particular, $\text{bisector}_{(F,G)}(a, b)$ is empty unless $\{b_i \mid i \in L_{+1}\} = \{b_{d+1}\}$ is a singleton. Since the coefficients of b are in ascending order, it follows that

$$\gamma + \delta \ = \ b_{d+1} \tag{2.14}$$

and $b_i = b_{d+1}$ for all $i \in L_{+1}$. This shows that (II.1) holds. Now the constraints of (2.9) translate into (2.13) as in the previous case, which implies (I.1). Additionally

$$\gamma - \delta \ \leq \ b_1 \ . \tag{2.15}$$

Adding (2.14) and (2.15) now yields

$$\gamma \ \leq \ \tfrac{1}{2}(b_1 + b_{d+1}) \ = \ \mu(b) \ ,$$

which, by (2.13), gives (II.2). Finally, (II.3) follows from the fact that the only possible value in $\{b_i \mid i \in L_0\}$ is γ.

The case where $L_{-1} \neq \emptyset$ and $L_{+1} = \emptyset$ is analogous.

Claim III: *Suppose that $L_{+1} \neq \emptyset \neq L_{-1}$. Then,* (2.9) *is feasible if, and only if we have*

$$\begin{aligned} &\{b_i \mid i \in L_{-1}\} = \{b_1\} \ , \qquad \{b_i \mid i \in L_{+1}\} = \{b_{d+1}\} \ , && \text{(III.1)} \\ &b_i \leq \mu(b) \ \textit{for} \ i \in L_- \ , \qquad b_i = \mu(b) \ \textit{for} \ i \in L_0 \ , \qquad b_i \geq \mu(b) \ \textit{for} \ i \in L_+ \ . && \text{(III.2)} \end{aligned}$$

Indeed, in this case the only candidate solution for (2.9) is $\gamma = \mu(b) = (b_{d+1} + b_1)/2$ and $\delta = (b_{d+1} - b_1)/2$. This is a solution or not depending only on whether Equations (III.1) and (III.2) are satisfied. $\square$

The following result gives the second direction of Theorem 2.4.6: $\mathcal{F}^d_{\text{bop}}$ is coarser than $\mathcal{F}^d_{\text{bis}}$.

Lemma 2.4.11. *Let $a, b, a', b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ be two pairs of points. If the bisected ordered partitions of $b - a$ and $b' - a'$ are not the same, then there is a pair of faces $F, G \in \mathcal{F}(\mathbb{B}^d)$ such that exactly one of* $\text{bisector}_{F,G}(a, b)$ *or* $\text{bisector}_{F,G}(a', b')$ *is empty.*

Proof. As before, we assume without loss of generality that $a = a' = 0$. We know that the bisected ordered partitions of b and b' are different. Our goal is to find a pair of faces (F, G) that lies in one and only one of the bisectors. We do this in three cases, depending of what is the difference between the bisected ordered partitions

Case I: *Suppose that*

$$\left\{ i \in [d+1] \,\middle|\, b_i = \max_j b_j \right\} \neq \left\{ i \in [d+1] \,\middle|\, b'_i = \max_j b'_j \right\} \text{ or}$$
$$\left\{ i \in [d+1] \,\middle|\, b_i = \min_j b_j \right\} \neq \left\{ i \in [d+1] \,\middle|\, b'_i = \min_j b'_j \right\} .$$

Without loss of generality that there is an i such that b_i is maximum and b'_i is not, or b_i is minimum and b'_i is not. Let F be the face with

$$F_+ = \left\{ i \in [d+1] \,\middle|\, b_i = \max_j b_j \right\} , \qquad F_- = \left\{ i \in [d+1] \,\middle|\, b_i = \min_j b_j \right\} ,$$

and let $G = -F$. This choice makes

$$L_{+1} = F_+ , \qquad L_{-1} = F_- , \qquad L_* = F_* , \qquad L_- = L_+ = L_0 = \emptyset .$$

The cell $\text{bisector}_{F,G}(a, b)$ is not empty by lemma 2.4.7, since the following is a solution for (2.9):

$$\begin{aligned}
\gamma &= \frac{1}{2}\left(\max_{i\in[d+1]} b_i + \min_{i\in[d+1]} b_i\right) , \\
\delta &= \frac{1}{2}\left(\max_{i\in[d+1]} a_i - \min_{i\in[d+1]} b_i\right) = \frac{1}{2}\,\text{dist}(a, b) .
\end{aligned}$$

However, $\text{bisector}_{(F,G)}(a', b')$ is empty: in order for it not to be empty we would need

$$\left\{ i \in [d+1] \,\middle|\, b'_i = \min_j b'_j \right\} \subset F_- , \qquad \left\{ i \in [d+1] \,\middle|\, b'_i = \max_j b'_j \right\} \subset F_+ .$$

Case II: *Suppose that b and b' have exactly the same maxima and minima but the ordered partitions of b and b' do not coincide.* That is, there is a pair of indices, $i, j \in [d+1] \setminus \{1, d+1\}$ such that $b_i \geq b_j$ but $b'_i < b'_j$.

We assume without loss of generality that 1 and $d+1$ are a minimum and a maximum, respectively, of both b and b'. Let F be the face with $F_+ = \{d+1\}$ and $F_- = \{1\}$. Let G be the face with $G_+ = \{i\}$, and $G_- = \{j\}$. Then, (2.8) gives us that

$$L_+ = \{d{+}1, j\} , \qquad L_- = \{1, i\} , \qquad L_* = [d{+}1]\backslash(L_+ \cup L_-) , \qquad L_{-1} = L_{+1} = L_0 = \emptyset .$$

Then, $\text{bisector}_{F,G}(a, b)$ is not empty since $\gamma = (b_i + b_j)/2$, and $\delta = \text{dist}(a, b)$ is a solution of (2.9).

However, the system for b' is unfeasible, since $b'_i < b'_j$. Therefore, $\text{bisector}_{F,G}(a', b')$ is empty.

Case III: *Suppose that b and b' have exactly the same maxima and minima and the same ordered partitions but the midvalue does not coincide.*

As before, we assume without loss of generality that 1 and $d+1$ are a minimum and a maximum, respectively, of both b and b'. Then, there is an index $i \in [d+1] \setminus \{1, d+1\}$ such that $\mu(b) \leq b_i$ but $\mu(b') > b'_i$ (or vice-versa, but that would give an equivalent case).

In this case, we let F and G be the faces with $F_+ = \{d+1\}$, $F_- = \{1\}$, $G_+ = \{1\}$ and $G_- = \{i\}$. These faces produce

$$L_+ = \{d{+}1, i\} \ , \qquad L_{-1} = \{1\} \ , \qquad L_* = [d{+}1] \setminus (L_+ \cup L_{-1}) \ , \qquad L_- = L_{+1} = L_0 = \emptyset \ .$$

Then, $\mathrm{bisector}_{F,G}(a,b)$ is not empty since $\gamma = \mu(b)$, $\delta = \mathrm{dist}(a,b)/2$ is a solution of (2.9) for b

However, the system for b' is unfeasible. This is because (2.9) specifies that $\gamma + \delta \geq b'_{d+1}$, and $\gamma - \delta = b'_1$. Adding them together and dividing by two we get $\gamma \geq (b'_1 + b'_{d+1})/2 = \mu(b')$. We also need by (2.9) and $i \in L_+$ that $\gamma \leq b'_i$. Then, $\mu(b) \leq b'_i$, which contradicts our assumption. $\square$

Corollary 2.4.12. *The tropical bisector of two points in general position has* $\Theta(d^4)$ *maximal cells.*

Proof. The upper bound is trivial, each maximal cell corresponds to a choice of a pair of facets (F,G) from the tropical ball. For the lower bound, assume without loss of generality that $a = 0$ and $b_1 < b_2 < \cdots < b_{d+1}$. Then, for each choice of $i, j, k, \ell \in \{1, \ldots, d+1\}$, all different and with $\max\{j, \ell\} < \min\{i, k\}$, let $F_+ = \{i\}$, $F_- = \{j\}$, $G_+ = \{k\}$, $G_- = \{\ell\}$. By Claim 1 in the proof of Lemma 2.4.10 the set $\mathrm{bisector}_{(F,G)}(a,b)$ is not empty. Since a and b are in general position and F and G are facets of $\mathbb{B}^d$, we have $\dim \mathrm{bisector}_{(F,G)}(a,b) = d-1$. There are $4\binom{d+1}{4}$ ways of choosing such $\{i,j,k,\ell\}$. $\square$

2.4.4 The structure of tropical Voronoi regions

A *polytrope* is an ordinary polytope which is also convex in the tropical sense (with respect to min and max simultaneously); cf. [36]. These are precisely the ordinary polytopes whose facets normals are roots of type A_d, i.e., $e_i - e_j$ for $i \neq j$; they generalize the "alcoved polytopes" of Lam and Postnikov [48]. Here we relax this notion by also calling a not necessarily bounded ordinary polyhedron a *polytrope* if its facets normals are roots of type A_d; this was called a "weighted digraph polyhedron" in [37].

The tropical unit ball $\mathbb{B}^d$ is a polytrope. But a more important example for us are the polytropes $Q = \bigcap_{a \in S}(a_+ F_a)$, where $F_a \in \mathbb{B}^d$ for each $a \in S$. Recall from Section 2.3 that in such a Q bisectors of subsets of S agree with affine subspaces. Thus:

Lemma 2.4.13. *For each polytrope Q as above and $a \in S$, the set $Q \cap \mathrm{Vor}_S(a)$ is the intersection of Q with ordinary affine halfspaces with facet normal $e_i - e_j - e_k + e_\ell$, where i and j are fixed.*

Proof. Let i and j be the coordinates maximized and minimized in F_a, respectively. For each $b \in S \setminus a$, the condition for x to be closer to a than to b is that $x_i - a_i - x_j + a_j \leq x_k - b_k - x_\ell + b_\ell$, where k and ℓ are the coordinates corresponding to F_b; cf. Proposition 2.4.1. $\square$

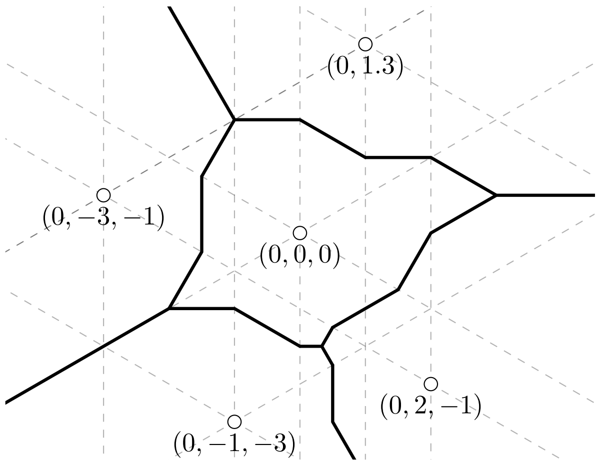

Figure 2.7: Tropical Voronoi diagram of five points in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$. The decomposition of Voronoi regions into semi-polytropes is shown by dashed lines

We call the intersection of a (possibly unbounded) polytrope with ordinary affine halfspaces with facet normal $e_i - e_j - e_k + e_\ell$, where i and j are fixed, a *semi-polytrope* of type (i, j). A semi-polytrope in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \cong \mathbb{R}^d$ has at most $2\binom{d+1}{2}$ facets, since there are at most $(d+1)d$ vectors $e_i - e_j - e_k + e_\ell$ for $k \neq \ell$ and fixed (i, j), plus the (at most) $(d+1)d$ facets of a polytrope.

A set $X \subset \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \cong \mathbb{R}^d$ is *star convex* with center c if for any point $x \in X$ the ordinary line segment $[c, x]$ is contained in X. Clearly any convex set is star convex, but the converse does not hold. Star convex sets are contractible. Despite the many differences to Euclidean Voronoi diagrams, the following result expresses a key similarity.

Theorem 2.4.14. *Let $S \subset \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ be a finite set in weak general position. Then each tropical Voronoi region of S is the star convex union of finitely many (possibly unbounded) semi-polytropes.*

Proof. That Voronoi regions for polyhedral norms are star-convex is a well-known fact (see [7, p. 133] or [54, p. 127]), which follows for example from Theorem 2.3.15. By Lemma 2.4.13, $\mathrm{Vor}_S(a)$ decomposes as finitely many semi-polytropes, by intersecting it with the individual polyhedra $Q = \bigcap_{a \in S}(a_+ F_a)$, for all choices of $\{F_a\}_{a \in S}$. □

Semi-polytropes are not necessarily tropically convex, and this entails that the regions of a tropical Voronoi diagram are not necessarily tropically convex either; cf. Figure 2.7 for an example. The tropical torus $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ is compactified by the tropical projective space $\mathbb{TP}^d$; the latter is the max-tropical convex hull of the $d+1$ max-tropical unit vectors

$$(0, -\infty, -\infty, \ldots, -\infty)\,,\ (-\infty, 0, -\infty, \ldots, -\infty)\,,\ \ldots\,,\ (-\infty, -\infty, \ldots, -\infty, 0)\ .$$

In this way, $\mathbb{TP}^d$ may be seen as an infinitely scaled tropical unit ball, which is a polytrope; cf. [37, §3.5]. Similarly for arbitrary (semi-)polytropes the line between bounded and unbounded is blurred in the compactification.

2.5 Computing tropical Voronoi diagrams

We will discuss several algorithms. Some of these methods are similar to their classical Euclidean counterparts, others rely on tailored data structures, which are based on Theorem 2.4.14. For their complexity analysis, we will consider the dimension as constant.

2.5.1 The planar case

There are several methods for computing Euclidean Voronoi diagrams in $\mathbb{R}^2$ with the optimal time complexity $O(n \log n)$ and linear space; cf. [19, §7.2]. This agrees with the situation for planar tropical convex hull computations; cf. [34, §5]. Chew and Drysdale [13] gave a divide-and-conquer algorithm with the same complexity for planar Voronoi diagrams with respect to arbitrary norms. Here we sketch a tropical analog of Fortune's beach line algorithm [25], which we discussed in Section 1.4.3; see also [68].

Suppose that we are given a set S of n sites in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$. In view of Theorem 2.4.14 the tropical Voronoi diagram of S gives rise to a planar graph where vertices are circumcenters of triples of points in S, edges are two point bisectors, and faces are Voronoi regions. We can make this planar embedding piecewise linear by subdividing each bisector into at most five segments; cf. Figure 2.4. The relevant data structure, as in the classical setting, is a doubly-connected edge list which requires $O(n)$ space; cf. [19, §2.2].

The beach line algorithm is based on a line sweep. The *tropical sweep line* at *time* t in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ is the set $L(t) = (0, t, 0) + \mathbb{R}(0, 0, 1) + \mathbb{R}\mathbb{1}$. Note that $L(t)$ is an ordinary line which is also tropically convex (with respect to min and max). For an arbitrary point x the set

$$P(x, t) \;=\; \left\{a \in \mathbb{R}^3/\mathbb{R}\mathbb{1} \;\middle|\; \operatorname{dist}(x, a) = \operatorname{dist}(x, L(t))\right\}$$

is the *parabola* spanned by x and $L(t)$; here $\operatorname{dist}(x, L(t)) = \min\{\operatorname{dist}(x, y) \mid y - (0, t, 0) \in \mathbb{R}(0, 0, 1) + \mathbb{R}\mathbb{1}\}$. This is a 1-dimensional polyhedral complex, which is homeomorphic with $L(t)$ via orthogonal projection, consisting of five segments.

A point $a = (a_1, a_2, a_3)$ is said to have been *visited* by the sweep line $L(t)$ if $a_2 - a_1 \leq t$. We will assume that our set S of sites is in general position and hence, in particular, the sweep line never visits two sites at the same time.

The *beach line* $B(t)$ of S at time t is formed by the points (b_1, b_2, b_3) which lie on a parabola $P(s, t)$ for a visited point $s \in$ such that $b_2 - b_1$ is maximal among all such points for a fixed value $b_3 - b_1$. That is, the beach line is formed by the right-most points on the parabolas spanned by the visited points and the sweep line; cf. Figure 2.8. So $B(t)$ is a union of parabolic arcs; it is easy to see that each parabola contributes at most two arcs to the beach line at any time. Like a single parabola also the beach line $B(t)$ is homeomorphic to $L(t)$ via orthogonal projection. In the portion of $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ left to $B(t)$ the tropical Voronoi diagram of S is known at time t.

Observation 2.5.1. *The beach line is a polygonal line with $O(n)$ segments.*

The actual algorithm works as in the classical case. We maintain a priority queue of *site events* (when the sweep line visits a site) and *circle events* (when there is a candidate for a new vertex of the tropical Voronoi diagram). The total number of events is linear in n. As in the classical case, it is possible to relax the condition on general position by means of symbolic perturbation.

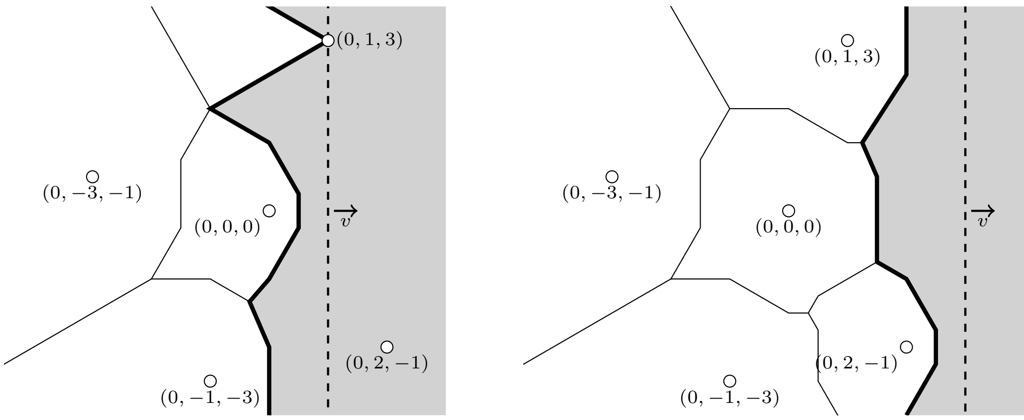

Figure 2.8: The beach line and the sweep line, at times $t = 1$ (left) and $t = 3$ (right). This is the tropical version of Figure 1.5.

Theorem 2.5.2. *The beach line algorithm computes a tropical Voronoi diagram of n sites in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ in $O(n \log n)$ time and $O(n)$ space.*

For the output we can choose, with the same complexity, between an abstract planar graph (encoding $\mathrm{Vor}(S)$ topologically) and its piecewise linear embedding resulting from Theorem 2.4.14.

2.5.2 Polytrope partitions

Let $S \subseteq \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ be a finite set of sites. From Theorem 2.4.14 we know that the tropical Voronoi diagram can be described in terms of (semi-)polytropes. For the definition and basic facts on polytropes, cf. Section 2.4.4 and [36]. The following takes inspiration from point location data structures, in particular the *trapezoidal map*; cf. [19, §6].

Definition 2.5.3. *A* polytrope partition *for S is a finite collection $\mathcal{C}$ of (perhaps unbounded) polytropes with disjoint interiors, covering $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$, such that:*

1. *Each facet-defining hyperplane of any cell in $\mathcal{C}$ lies in the hyperplane arrangement $S + A_d$.*
2. *For each cell P in $\mathcal{C}$ and site $a \in S$ the restricted Voronoi region $\mathrm{Vor}_S(a) \cap P$ is contained in a maximal cone $a + F$ of $a + \mathcal{F}(\mathbb{B}^d)$.*

A valid labeling *for $\mathcal{C}$ assigns to each cell $P \in \mathcal{C}$ a relation $\mathcal{L}_{\mathcal{C}}(P) \subseteq S \times \mathcal{F}(\mathbb{B}^d)$ between S and $\mathcal{F}(\mathbb{B}^d)$ containing*

$$\left\{(a, F) \in S \times \mathcal{F}(\mathbb{B}^d) \,\middle|\, (a + F) \cap P \cap \mathrm{Vor}_S(a) \neq \emptyset\right\}.$$

The labeling does not need to assign every site and facet, and it may have more than one site assigned to the same facet. But, by definition, every site will be assigned to at most one facet.

The idea behind this definition is that in each cell we restrict the number of sites we need to consider in order to study the Voronoi diagram in that cell, plus the distances to the relevant sites become linear.

Observation 2.5.4. *Let P be a cell in a polytrope partition $\mathcal{C}$ for S. Then for all $x \in P$ we have*

$$\operatorname{dist}(x, S) \;=\; \min_{a \in S} \lambda_{F_a}(x - a) \tag{2.16}$$

where λ_{F_a} is the linear function defined by restricting the distance to a on some maximal cone $a + F_a$ of $a + \mathcal{F}(\mathbb{B}^d)$ which contains $\operatorname{Vor}_S(a) \cap P$. Thus computing the restriction of $\operatorname{Vor}(S)$ to the polytrope P amounts to finding the regions of linearity of the tropical polynomial $\min_{a \in S} \lambda_{F_a}(x)$. The latter can be obtained via an ordinary dual convex hull computation. The facet F_a is exactly the facet matched to a in $\mathcal{L}_{\mathcal{C}}(P)$.

Example 2.5.5. *The braid arrangement A_d consists of the $\binom{d+1}{2}$ ordinary hyperplanes $\{x \mid x_i = x_j\}$, where $i \neq j$. This gives rise to the* standard polytrope partition $S + A_d$, *which is finer than any other polytrope partition for S; Figure 2.7 shows an example for $d = 2$. This construction occurs in planar tropical convex hull algorithms; cf. [34, Figure 3].*

For the labeling of cell P, we can just classify the sites $a \in S$ by which of the cones $F\mathcal{F}(\mathbb{B}^d)$ covers P, and assign to each facet the site associated with that facet closest to P, measured by the linear function with that normal vector. This will be the standard way we associate labels to each cell.

The standard polytrope partition is the finest (and most complex) possible polytrope partition. Coarser partitions will yield better algorithms.

For points in weak general position no valid labeling of any cell in a polytrope partition needs to use the same facet F twice, hence the labeling becomes a partial matching:

Lemma 2.5.6. *Let $\mathcal{C}$ be a polytrope partition for S. If S is in weak general position then there is a valid labeling of $\mathcal{C}$. Moreover, for d considered constant, a labeling of each polytropal cell has constant size, and it can be computed in $O(n)$ time.*

Proof. Suppose that a valid labeling does not exist. Then there are sites $a, b \in S$ and a maximal cone $F \in \mathcal{F}(\mathbb{B}^d)$ such that the sets $(a+F) \cap P \cap \operatorname{Vor}_S(s)$ and $(b+F) \cap P \cap \operatorname{Vor}_S(t)$ both are non-empty.

With the notation of (2.16) we have $\lambda_{F_a} = \lambda_{F_b}$; and we shortly write λ. Since the sites are in weak general position, we may assume that $\lambda(b) > \lambda(a)$. Picking $y \in (b + F) \cap P \cap \operatorname{Vor}_S(b)$ yields

$$\operatorname{dist}(y, b) \;\geq\; \lambda(b) \;>\; \lambda(a) \;\geq\; \operatorname{dist}(y, S) \;,$$

where the last inequality follows from (2.16). The resulting inequality $\operatorname{dist}(y, b) > \operatorname{dist}(y, S)$ implies that $y \notin \operatorname{Vor}_S(b)$, which is a contradiction. Hence a valid labeling does exist.

To compute such labeling, we iterate through all the sites. For each site a, the candidate facet of $F_a \in \mathcal{F}(\mathbb{B}^d)$ is known by definition of the polytrope partition. To check if (a, F_a) is a labeling candidate, we need to determine if $(a + F_a) \cap P \cap \operatorname{Vor}_S(a)$ is empty or not. This amounts to solving a linear program that has constant size (as d is a constant). It follows that the entire labeling can be computed in $O(n)$ time. ☐

We aim at a first algorithm for computing a tropical Voronoi diagram in arbitrary dimension. This will employ the standard polytrope partition from Example 2.5.5.

Lemma 2.5.7. *If S is in general position and has size n then the standard polytrope partition has*

$$(d+1)^{d-1}n^d + O(n^{d-1})$$

maximal cells, if we consider d a fixed constant.

Proof. Pick a generic direction $v \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. The cells of the polytrope partition that are bounded in the direction of v are in a one-to-one correspondence with the vertices of the arrangement, by associating each polytrope with the optimum of the linear program maximizing $v^T x$.

We compute the number of vertices in the arrangement in the following way. First we pick $d+1$ directions out of the $d(d+1)$ potential normals. Then, we choose which of the n choices for each direction will be the plane we are intersecting. The choice of normals require them to form a base. Since these normals correspond to the edges in the directed complete $(d+1)$ graph, and the bases correspond to spanning trees, we can count them by Cayley's formula. Finally, intersection of hyperplans centered at the same site have to be substracted. Thus, the exact number of vertices in the arrangement is

$$(d+1)^{d-1}n^d - n\left((d+1)^{d-1} - 1\right).$$

Now, it suffices to show that the number of unbounded cells is in $O(n^{d-1})$. The unbounded cells intersect a hyperplane, H, normal to v that is far enough in the v direction. The cells intersecting H are the same as the cells in the restricted hyperplane arrangement, which is a $(d-1)$-dimensional arrangement with $N = \binom{d+1}{2}n$ hyperplanes. The number of such cells is known to be in $O\left(N^{d-1}\right)$, which agrees with $O\left(n^{d-1}\right)$ as N depends linearly on n. ☐

Remark 2.5.8 (Standard polytrope partition algorithm)**.** *This directly yields a first algorithm for computing a tropical Voronoi diagram of n sites in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ in $O(n^{d+1})$ time, as follows: First, we sort S along each of the $\binom{(d+1)}{2}$ directions $e_i + e_j$, in $O(n \log n)$ time. As in the proof of Lemma 2.5.7 we pick a generic direction $v \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. We can compute the vertices of the hyperplane arrangement $A_d + S$ in time $O(n^d)$ by enumerating all d-sets of independent directions, which can be derived from the oriented spanning tree of K_{d+1}, in constant time. For each of the d directions we choose an index $i \in [n]$.*

Next we perturb each such vertex p by a small multiple of $-v$, and we collect the intersection of bands of contiguous parallel hyperplanes of $A_d + S$ that contain the perturbed point. This can be done in time $O(\log n)$ for each direction. In this way, we find those cells which are bounded in the direction of this particular v in linear output-dependent time. We repeat the same procedure for a set of directions $v_1, \ldots, v_{d+1}$ which positively span the entire space $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. Each polytropal cell will be bounded in at least one of these directions, and thus their enumeration is still in $O(n^d)$ for d fixed.

Then, for each polytrope P, we compute a corresponding labeling in time $O(n)$, by Lemma 2.5.6. Therefore, we can compute the standard polytrope partition, including labels, in time $O(n^{d+1})$ for fixed d. The tropical Voronoi diagram in each cell is an ordinary dual convex hull problem of constant size. This computation splits each polytrope in the partition into semi-polytropes. The convex hull problem can be solved in constant time, and hence this algorithm takes $O(n^{d+1})$ time, if d is considered a fixed constant.

Question 2.5.9. *In the plane $\mathbb{R}^3/\mathbb{R}\mathbb{1}$, we believe that ideas similar to the "trapezoidal maps" used in point location, cf. [19, §6.1], should yield polytrope partitions of linear size but we did not work out the details. More generally: Is there a polytrope partition of complexity better than $\Theta(n^d)$ in arbitrary dimension? One could hope for something in $O(n^{d/2})$, which is the worst-case complexity of Euclidean Voronoi diagrams.*

2.5.3 An $O(n^d \log n)$ randomized incremental algorithm in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$

We can improve the algorithm from Remark 2.5.8 by constructing a polytrope partition incrementally. The idea is to update an existing polytrope partition by including a new point and to employ randomization to improve the efficiency. Moreover, we will also produce a coarser polytrope partition than the standard one, but only by a constant in d.

A key ingredient is a new data structure that we call a *polytrope tree*. Throughout we assume that the set S of n sites forms a subset of $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ in general position. We fix the polytrope partition $\mathcal{C} := S + \mathcal{F}(\mathbb{B}^d)$, which is coarser than the standard polytrope partition but only by a factor which is constant in d; cf. Example 2.5.5.

Definition 2.5.10. *A* polytrope tree *for S is a (rooted) tree T such that*

1. *for each leaf ℓ there is a polytropal cell $P(\ell)$ of $\mathcal{C}$;*
2. *for each interior node i there is a site $a(i) \in S$ and a polytrope $P(i)$.*

These satisfy the following consistency conditions:

- *for the root node r of T we have $P(r) = \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$, which may be seen as an unbounded polytrope;*
- *the map $\ell \mapsto P(\ell)$ is a bijection between the leaves of T and the polytropes in $\mathcal{C}$;*
- *the map $i \mapsto a(i)$ is a surjection from the interior nodes onto the set S;*
- *if i is an interior node with children $c_1, \ldots, c_k$, then $P(c_1), \ldots, P(c_k)$ form the maximal cells of $(a(i) + \mathcal{F}(\mathbb{B}^d)) \cap P(i)$.*

It is easy to construct a polytrope tree for S, and its purpose is to speed up the computation of a valid labeling. This will reduce the algorithmic complexity from $O(n^{d+1})$ to $O(n^d \log n)$.

For the incremental update to insert a new site $b \notin S$ we maintain a stack Σ of unvisited nodes in a given polytrope tree for Σ and process it as follows:

- the stack Σ is initialized with the root node r;
- we remove the top node q from the stack Σ unless it is empty;
- if q is an interior node such that $P(q)$ intersects more than one maximal cone of $s + \mathcal{F}(\mathbb{B}^d)$, then we push the children of q onto the stack Σ;

- if q is a leaf such that $P(\ell)$ intersects more than one maximal cone of $p+\mathcal{F}(\mathbb{B}^d)$, then we create the intersections of $P(\ell)$ with $s+\mathcal{F}(\mathbb{B}^d)$ as new leaves, which now become children of q, and we set $a(q) \leftarrow b$.

Note that an interior node q with $P(q)$ contained in a unique maximal cone of $a+\mathcal{F}(\mathbb{B}^d)$ is kept unchanged, and its children will not be visited. The following is the essential part of the complexity analysis.

Proposition 2.5.11. *Let T be a polytrope tree created in the way explained above, where the n sites in S are processed in uniformly random order. Then the expected height of T is of order $O(\log n)$, if d is considered a fixed constant.*

Proof. Let P be a polytrope in the polytrope partition $\mathcal{C}(S)$. For each ordering $\pi : [n] \to S$ of S we have a polytrope tree $T(S,\pi)$ with P as a leaf. By induction on n we will show:

$$E[h_{T(S,\pi)}(P)] \ \le \ d(d+1)\sum_{i=1}^{n}\frac{1}{i} \quad \in \ O(\log n) \ , \tag{2.17}$$

where the expectation $E[\cdot]$ is taken uniformly over all $n!$ orderings of S, and $h_{T(S,\pi)}(P)$ is the depth of the leaf of P in $T(S,\pi)$.

We proceed by backwards analysis. Let $S' \subset S$ be the subset of sites that lie in some facet-defining hyperplane of P. Since P has at most $d(d+1)$ facets and (by general position) their corresponding hyperplanes contain each exactly one point of S, we have $|S'| \le d(d+1)$. Thus, the probability that the height $h_T(P)$ increases in the last insertion is at most $d(d+1)/n$. Since the increase is by exactly one, we have

$$E[h_{T(S,\pi)}(P)] \ \le \ E[h_{T(S\setminus\pi(n),\pi_{[n-1]})}(P')] + \frac{d(d+1)}{n} \ ,$$

where P' is the polytrope containing P in the polytrope partition before the last insertion. By induction hypothesis

$$E[h_{T(S\setminus\pi(n),\pi_{[n-1]})}(P')] \ \le \ d(d+1)\sum_{i=1}^{n-1}\frac{1}{i} \ .$$

The last two formulas give Equation (2.17). □

Corollary 2.5.12. *The above method constructs a polytrope tree for the polytrope partition $S+\mathcal{F}(\mathbb{B}^d)$ in expected time $O(n^d \log n)$ and space $O(n^d)$, for d constant.*

Proof. The algorithm that inserts a new site a into the tree only visits nodes that are above some leaf requiring an update. For each such leaf ℓ the polytrope $P(\ell)$ intersects one of the $d(d+1)$ hyperplanes in $a+A_d$. This implies that there are $O(n^{d-1})$ of them. Since the expected depth of every leaf is $O(\log n)$ it requires expected time $O(n^{d-1}\log n)$ for inserting a. Hence the total complexity for n sites amounts to $O(n^d \log n)$. □

In order to compute the tropical Voronoi diagram, we also need to compute the labeling of this polytrope partition. The naive way is to compute the labeling for each leaf as we did in Remark 2.5.8.

A slight improvement is to compute the labeling during the depth first search (DFS) exploration of the tree at each insertion of a new site. But in this way, even if an interior node is completely contained in only one cone of the fan $a + \mathcal{F}(\mathbb{B}^d)$, we need to descend to its subtree in order to update the labels. This would slow the algorithm down to $\Theta(n^{d+1})$ because each insertion will have to iterate through all the leaves.

A better way is to compute the labeling lazily. To this end we equip each interior node i with a partial labeling $\mathcal{L}_{\mathcal{C}}(i)$. With each new insertion, we proceed as we just explained, but we do not cascade down the label updates. Only once all sites in S have been inserted we cascade the partial labels, updating them in DFS order. This takes $O(n^d \log n)$ time to compute the polytropes and the lazy labellings, plus $O(n^d)$ time to cascade the lazy labellings down in the tree, for a total time complexity of $O(n^d \log n)$ time. This gives our final result.

Theorem 2.5.13. *There is a randomized incremental algorithm for computing tropical Voronoi diagrams of n sites in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ in general position with expected time complexity $O(n^d \log n)$ and space complexity $O(n^d)$, for d constant.*

2.5.4 Coarser partitions and empty polytropes

In the previous section, when discussing how to construct TVDs in $O(n^d \log n)$ time in fixed dimension d, we did not address how to compute the intersection of each polytrope in the tree with the fan $p + \mathcal{F}(\mathbb{B}^d)$ centered at the point being added. This is not important for its asymptotic complexity, since both the polytrope and the fan have complexity bounded by d independently of n, so this step is $O(1)$ in fixed dimension no matter how we do it.

Yet, in this section we show a way to do this step in polynomial time in d using graph algorithms (instead of the potentially more expensive linear programming), and we give an improved criterion to split polytropes during the execution of the algorithm in Theorem 2.5.13. This is related to the parametrized shortest paths construction in [40].

Lemma 2.5.14. *Let*

$$P = \{x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} : x_i - x_j \leq B_{ij} \ \forall i, j \in [d+1]\}$$

be a polytrope defined by a matrix $B \in \mathcal{M}_{(d+1)\times(d+1)}(\mathbb{R} \cup \{\infty\})$. Let $F \in \mathcal{F}\mathbb{B}^d$ be some cone. There is an algorithm that decides if $(p + F) \cap P = \emptyset$ in time $O(d^3)$. If B is a shortest path matrix (i.e., $B = B^*$, and all hyperplanes touch the polytrope at some point), then there is an algorithm that decides if $F \cap P = \emptyset$ in time $O(d^2)$.

Proof. Since the cone $p+F$ is a polytrope and the intersection of polytropes is a polytrope, the problem reduces to determining if a polytrope is empty or not.

In particular, let F_+, F_- be the indices of the maximum and minimum in F (usually these are sets, but since F is a facet, we consider them elements of $[d + 1]$). Then, the polytrope $B \cap (p + F)$ is defined by the inequalities:

$$\begin{cases} x_i - x_j \leq \min(B_{ij}, p_i - p_j) & \text{if } i = F_+ \text{ or } j = F_- \\ x_i - x_j \leq p_i - p_j & \text{otherwise.} \end{cases}$$

It is known [35, Chapter 3] that a polytrope defined by a matrix M is empty if and only if M does not accept a shortest path form M^*, that is, if the directed weighted graph defined by M on $[d+1]$ has no negative cycle. We can compute M^* in time $O(d^3)$ by Floyd-Warshall algorithm.

For the $O(d^2)$ improvement, observe that when we intersect P with the new cone we modify exactly one row and one column of B. If we assume that $B = B^*$ is in shortest paths form, it suffices to perform two iterations of Floyd-Warshall, on the vertices F_+ and F_-. This is because we could run Floyd-Warshall us usual but considering F_+ and F_- last. Since B is already in shortest paths form, the first $d-2$ iterations will not change the matrix. We can skip them, and just run the last two iterations, which take $O(d^2)$ time. ☐

Remark 2.5.15. *The algorithm used in Theorem 2.5.13 creates always the same polytrope map, namely $S + \mathcal{F}(\mathbb{B}^d)$, independently of the order in which the points of S are inserted. A better idea would be to divide only the polytropes that intersect the Voronoi region of the point p being inserted, since for the others we do not need to know in what cone of $p + \mathcal{F}(\mathbb{B}^d)$ they are (and hence we do not care if they intersect more than one cone). The following lemma shows how to detect what polytropes need to be divided:*

Lemma 2.5.16. *Let S be points in weak general position, P be a polytrope of a polytrope map $\mathcal{C}$ of S, let $\mathcal{L}$ be a labeling of P, and let $p \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \setminus S$ be some new point. Let $F \in \mathcal{F}(\mathbb{B}^d)$ be a cone. Then, there is an algorithm that decides if $P \cap (p+F) \cap \mathrm{Vor}_{S\cup\{p\}}(p) = \emptyset$ in time $O(d^4)$. The facets of the intersection $P \cap \mathrm{Vor}_{S\cup\{p\}}(p) \cap p + \mathcal{F}(\mathbb{B}^d)$ can be computed then in $O(d^6)$.*

Proof. The problem is to determine if there is some positive real number $R > 0$ an $x \in P \cap p + F$ such that $p \in \mathbb{B}^d(x, R)$ but $S \not\subseteq \mathbb{B}^d(x, R)$

When R is fixed, this problem is just an emptiness test on a polytrope. The equations are the facets of P and $p + F$, plus the inequalities $\mathrm{dist}(x, p_i) = \mathrm{dist}_{\mathcal{L}(p_i)}(x, p_i) \geq R$ and $\mathrm{dist}(x, p) = \mathrm{dist}_F(x, p) \leq R$. All these linear inequalities have normal vectors in A_n, therefore the feasible region is a polytrope

In the graph theoretic interpretation of the emptiness test for polytropes, we have that the polytrope is empty if and only if there is a negative cycle in the graph defined by matrix B.

When we consider the inequalities of P, plus the conditions on R, we have the same problem with a parametrized graph, where some weights are a tropical polynomial. These polynomials are Laurent polynomials of the form $x_i - x_j \leq \min(B_{ij}, C_{ij} - R)$ for some matrix $C \in \mathcal{M}_{(d+1)\times(d+1)}(\mathbb{R} \cup \{\infty\})$.

If this graph has a negative cycle (for some R), then it has a negative cycle of length at most $(d+1)$. Therefore, we can apply Floyd-Warshall as usual on the parametrized graph, to look for negative k-cycles for some k. After $d+1$ iterations, the length of the smallest cycle is a tropical $(d+1)$-polynomial on $-R$. We can determine when this polynomial is negative in $O(d+1)$ time, and we can compute the polynomial itself in $O(d^3)$ operations on tropical polynomials, each taking $O(d)$ time and $O(d)$ space.

☐

Even though this algorithm is, for each individual polytrope, slower than the algorithm described in Theorem 2.5.13 for fixed d, it produces a coarser polytrope tree/map, because some polytropes that intersect the new cone may be too far from that Voronoi region, and do not need splitting. However, the global behaviour is more difficult to analyze and we do not know if the complexity of this version has significatively better asymptotics. It is at least as efficient, and we expect it runs much faster in practice.

Chapter 3

Diameter of pure complexes

This chapter is from the article "Randomized construction of simplicial complexes with large diameter" [15], by Francisco Criado and Andrew Newman, published in Discrete & Computational Geometry in 2020. This article is available at `https://doi.org/10.1007/s00454-020-00248-2`. It is licensed under a Creative Commons Attribution 4.0 International License.

3.1 Introduction

Recall from Section 1.1 that, given a pure simplicial complex C, one may define the *dual graph* $\mathrm{G}(C)$ as the graph whose vertices are the facets of C and whose edges are pairs of facets intersecting in a ridge. We say that the combinatorial diameter of C is the graph diameter of $\mathrm{G}(C)$. Of course, in general this may be infinite as $\mathrm{G}(C)$ need not be connected. Thus, here we consider the case where C is *strongly connected*, that is exactly the case that $\mathrm{G}(C)$ is connected.

The best known situation in which the combinatorial diameter is considered is in the case that C is a simplicial polytope. In this situation there is the now-disproved [63] Hirsch conjecture which claimed that if C is a simplicial polytope of dimension d with n vertices, then the diameter of C is at most $n-d$. While this is now known to be false, the polynomial Hirsch conjecture remains open. The polynomial Hirsch conjecture asserts that the combinatorial diameter of every simplicial d-polytope on n vertices is bounded by a polynomial in n and d.

On the other hand, the question can also be considered for other classes of simplicial complexes as a purely combinatorial question. Following the notation of [16], we define $H_s(n,d)$ to be the largest combinatorial diameter of any strongly connected, pure $(d-1)$-dimensional simplicial complex on n vertices. Recall that d is not the dimension of the complex, rather $d-1$. This is to respect the notation for polyhedra, in which d refers to the dimension of the ambient space. In general, one could think of d as the number of vertices in a facet.

In [16], Criado and Santos prove the following:

Theorem 3.1.1 (Theorem 1.2 of [16])**.** *For every $d \in \mathbb{N}$ there are infinitely many $n \in \mathbb{N}$*

such that

$$H_s(n,d) \geq \frac{n^{d-1}}{(d+2)^{d-1}} - 3.$$

Combined with a trivial upper bound of $n^{d-1}/((d-1)(d-1)!)$, which appears as Corollary 2.7 of [45], this shows that $H_s(n,d) = \Theta_d(n^{d-1})$. This result of [16] gives the previously best known upper and lower bounds on $H_s(n,d)$, but as d tends to infinity the ratio between the upper bound and the lower bound grows like $e^{\Theta(d)}$. Here we take a step further toward establishing the true value of $H_s(n,d)$ in decreasing this ratio to $\Theta(d^2)$. Specifically, we prove the following lower bound:

Theorem 3.1.2. *Fix $d \geq 3$, then $H_s(n,d)$ satisfies*

$$(1 - o_n(1))\frac{1}{4ed^2} \leq \frac{H_s(n,d)d!}{n^{d-1}} \leq \frac{d}{d-1}.$$

The upper bound in Theorem 3.1.2 is the trivial one, which we get by double counting.

Our proof takes a different approach than the deterministic construction in [16]. We instead use the probabilistic method in a way similar to the main result of [59].

We make the comparison between what we do here and what is done in [59] more precise when we outline the proof below. At a basic level, there is a deterministic step and then a probabilistic step. We start with the dimension $(d-1) \geq 2$ and a positive integer L that we want to realize as the combinatorial diameter of our construction. In the deterministic step we build a complex on $\Theta(L)$ vertices that has diameter L. In the probabilistic step we take a quotient of the complex that preserves the diameter, drops the number of vertices to $\Theta(L^{1/(d-1)})$, and remains a simplicial complex.

This approach also works for the class of pseudomanifolds, which is also considered in [16]. A pseudomanifold is a pure dimensional and strongly connected simplicial complex so that every ridge is contained in exactly two facets. We denote by $H_{pm}(n,d)$ the maximum diameter of all $(d-1)$-dimensional pseudomanifolds on n vertices. A result of [16] shows that $H_{pm}(n,d) = \Theta_d(n^{d-1})$, but again the ratio between the upper bound and the lower bound is $e^{\Theta(d)}$. Here we improve this ratio to $\Theta(d^3)$:

Theorem 3.1.3. *Fix $d \geq 3$, then $H_{pm}(n,d)$ satisfies*

$$(1 - o_n(1))\frac{1}{4ed^4} \leq \frac{H_{pm}(n,d)d!}{n^{d-1}} \leq \frac{6}{(d+1)}.$$

In this case we slightly improve on the upper bound too by using the fact that $\mathrm{G}(C)$ is d-regular when C is a $(d-1)$-pseudomanifold.

3.2 Proof of main result

The approach will be as in [59], the goal of which is to construct simplicial complexes on few vertices with large torsion groups in homology. The main result of [59] shows that for any finite abelian group G and dimension d, there exists a simplicial complex on $O_d(\log^{1/d}|G|)$ vertices which realizes G as its top cohomology group. The construction in [59] is partially deterministic and partially probabilistic. The deterministic piece constructs a simplicial complex C on $O_d(\log|G|)$ vertices that realizes G as the top cohomology group. The probabilistic piece is to use the Lovász Local Lemma to color the vertices of C in a way that allows us to take a quotient of the complex by the coloring to obtain a simplicial complex on the right number of vertices, but without changing the top cohomology group. This technique was further refined in [60]. Both in [59, 60] and here, once we have found a good coloring and a good initial construction the quotient is taken according the following, combinatorial definition.

Definition 3.2.1. *([59]) If C is a simplicial complex with a coloring f of $V(C)$ we define the* pattern *of a face to be the multiset of colors on its vertices. If f is a proper coloring, in the sense that no two vertices connected by an edge receive the same color, we define the* pattern complex *C/f to be the simplicial complex on the set of colors of f so that a subset P of the colors of f is a face of C/f if and only if there is a face of C with P as its pattern.*

We also introduce the term *pattern classes of f on C* to refer to equivalence classes under the equivalence relation given by faces of C with the same pattern under f.

Our initial construction, that is the deterministic step, for Theorem 3.1.2 is quite simple. For dimension $d-1$ fixed, we define the *straight $(d-1)$-corridor on N vertices* to be the pure complex $SC(N,d)$ on $[N] := \{1,2,...,N\}$ where the facets are given by $\{1,...,d\},\{2,...,d+1\},\{3,...,d+2\},...,\{N-d+1,...,N\}$. Clearly $SC(N,d)$ has N vertices and its dual graph is a path of length $N-d$, so the diameter of $SC(N,d)$ is $N-d$. For the probabilistic step we want to color the vertices by a coloring f with $O_d(N^{1/(d-1)})$ colors so that $SC(N,d)/f$ still has diameter $N-d$. In fact we take a quotient so that the resulting complex will still have the same dual graph as $SC(N,d)$. The number of vertices will decrease drastically after this identification.

The rule that the coloring f should satisfy is that it should be a proper coloring and it should assign no pair of ridges the same pattern. This rule for coloring vertices is exactly the same as the rule for the result in [59].

If C is a simplicial complex and f is a coloring that colors every ridge with a unique pattern then C and C/f will have the same diameter; in fact they will have the same dual graph. Also under this condition if C is a pseudomanifold then so is C/f. These facts can be proved by comparing the $(d-1)$st boundary matrix of C and of C/f.

Recall that the i-dimensional boundary matrix ∂_i over $\mathbb{Z}/2\mathbb{Z}$ of a simplicial complex C is a matrix over $\mathbb{Z}/2\mathbb{Z}$ with columns indexed by the i-dimensional faces of C, rows indexed by $(i-1)$-dimensional faces of C, and the (σ,τ)-entry equal to 1 if and only if $\sigma \subseteq \tau$.

Lemma 3.2.2. *If C is a $(d-1)-$simplicial complex and f is a proper coloring of C so that every ridge of C has a unique pattern then the top-dimensional boundary matrix of C over $\mathbb{Z}/2\mathbb{Z}$ is the same as the top-dimensional boundary matrix of C/f. In particular if C is a pseudomanifold then so is C/f, and C and C/f have the same dual graph.*

Proof. We have that $\phi : V(C) \to V(C/f)$ by sending v to $f(v)$ induces a simplicial map and by the assumption on f it is injective on the set of ridges and on the set of facets. Moreover it is clear that for any ridge σ and any facet τ one has $\sigma \subseteq \tau$ if and only if $\phi(\sigma) \subseteq \phi(\tau)$. From this it is immediate that over $\mathbb{Z}/2\mathbb{Z}$, both C and C/f have the same top-dimensional boundary matrix.

A $(d-1)$-pseudomanifold is a $(d-1)$-complex C so that then every row of ∂_{d-1} of C has exactly 2 non-zero entries, so if C is a pseudomanifold then so is C/f. The dual graph of C may be read off of ∂_{d-1} as the off-diagonal entries of $\partial_{d-1}^T \partial_{d-1}$ match those of the adjacency matrix of $\mathrm{G}(C)$, so C and C/f have the same dual graph. $\square$

For the probabilistic step of our proof we want to properly color the vertices in $SC(N,d)$ using $O_d(N^{1/(d-1)})$ colors so that no pair of ridges receives the same pattern. Then the pattern complex will still have diameter $N-d$.

As in [59] and [60] the coloring is done in steps. In particular, there will be two steps to the coloring process. The main reason for this is that the first step allows us to handle pairs of ridges which intersect one another while the second step will handle ridges that do not intersect one another.

For fixed dimension $d-1$, the two step approach first colors $SC(N,d)$ by a coloring f with $O_d(1)$ colors so that no vertices that are at distance at most two from one another in the 1-skeleton of $SC(N,d)$ receive the same color. This condition implies that there are no pair of intersecting ridges receiving the same pattern. In the second step, one uses the Lovász Local Lemma to color the vertices of $SC(N,d)$ by a coloring g having $O_d(N^{d-1})$ colors so that no pair of disjoint ridges receive the same pattern. The product of the two colorings (f,g), assigning the pair $(f(v),g(v))$ to each vertex v, is then the final coloring that we use.

There is another advantage to this two step approach as well. Namely, we may regard the second coloring as a refinement of the first. Indeed if two disjoint ridges receive different patterns by f, then it doesn't matter what happens on them under the coloring by g as they will still receive different patterns. What this allows for is that if we use *more* colors for f, though still only a number depending on d, we can save on the number of colors we need for g in a way that is a net reduction in the number of colors in (f,g).

It is not too hard to check that the 1-skeleton of $SC(N,d)$ has maximum vertex degree $2(d-1)$, and so by a greedy coloring with at most $[2(d-1)]^2+1$ colors we may color the vertices so that no pair at distance two receives the same color. This can be refined using a coloring g obtained randomly using the Lovász Local Lemma, but this turns out to give a worse lower bound on $H_s(n,d)$ than the bound in [16], though it is still $\Theta_d(n^{d-1})$.

Here instead, it is better to take a random coloring for f which we describe below. This will ultimately allow for fewer colors for g and overall.

3.2.1 The first coloring

The purpose of this section is to prove the following:

Proposition 3.2.3. *Let $d \geq 3$, $c_1 > 6(d-1)$, and $\varepsilon > 0$ fixed. There is a coloring of $SC(N,d)$ with c_1 colors such that all pattern classes of ridges have at most size $(1+\varepsilon)N(d-1) \Big/ \binom{c_1}{d-1}$ and no pair of intersecting ridges have the same pattern for N large enough.*

Our proof makes use of the following result on Markov chains:

Theorem 3.2.4. *[Theorem 1.10.2 in [61]] Let $(X_n)_{n\geq 0}$ be an irreducible Markov chain on a finite set of states S, that is (X_n) is a Markov chain on S in which every state can reach any other with nonzero probability in an arbitrary number of steps. Then,*

$$\Pr\left(\frac{V_i(n)}{n} \to \frac{1}{m_i} \text{ as } n \to \infty\right) = 1\ \forall i \in S,$$

where m_i is the expected number steps to go from state i to itself, and $V_i(n)$ is the number of visits to i before n steps.

Proof of Proposition 3.2.3. We assign a sequence of colors to the vertices $v_1, \ldots, v_N$ of $SC(N,d)$ by a greedy randomized approach. To each vertex we assign a uniformly random color of $[c_1]$ that was not used in the last $2(d-1)$ vertices. This is always possible, as $c_1 > 6(d-1) > 2(d-1)$.

If two intersecting ridges share a pattern, in particular it means that two vertices, one from each ridge but not from the intersection have the same color. Since the facets of $SC(N,d)$ are consecutive d-sequences of vertices, two vertices from two intersecting ridges have to be at most $2(d-1)$ units apart with respect to the ordering on the vertices. This proves that the random procedure will always produce a coloring so that no pair of intersecting ridges have the same pattern.

It remains to prove that the pattern classes have the claimed bounded size. Observe that the randomized coloring procedure defined above defines a Markov Chain on the set:

$$S = \left\{x \in [c_1]^{2(d-1)} : x_i \neq x_j \forall i,j \in [2(d-1)], i \neq j\right\},$$

coming from the tuple of colors on $2(d-1)$ consecutive vertices.

In this Markov chain, we can get from any state $x \in S$ to any other $y \in S$ by first replacing the colors of x with colors not in x or y, which is possible because $c_1 > 6(d-1)$ and each state has $2(d-1)$ colors, and then replacing these new colors with the colors of y. This takes $4(d-1)$ steps independently of the initial and final states. Therefore, the Markov chain is irreducible.

By Theorem 3.2.4, as N grows large enough, the proportion of visits to each state approximates the expected value arbitrarily, with probability 1. By symmetries via permutations of colors, this proportion of visits has to be the same for each state, that is,

$$\frac{1}{\binom{c_1}{2(d-1)}(2(d-1))!} = \frac{(c_1-2(d-1))!}{c_1!}.$$

This means that for $\varepsilon > 0$ and N large enough there is a coloring f of $SC(N,d)$ by c_1 colors so that for every tuple of $2(d-1)$ colors of $[c_1]$, the proportion of tuples of $2(d-1)$ consecutive vertices colored by that tuple of colors is at most $(1+\varepsilon)(c_1-2(d-1))!/c_1!$.

Now each ridge σ of $SC(N,d)$ is given by $\sigma_{i,j} := \{i+1, ..., i+d\} \setminus \{i+j\}$ for some $i \in \{0, ..., N-d\}$ and $j \in \{1, .., d-1\}$ except for the ridge $\{1, ..., d-1\}$. For each ridge $\sigma_{i,j}$ in $SC(N,d)$ with $i \geq d-2$ we may associate the set of $2(d-1)$ consecutive vertices ending at $i+d$. This canonically associates almost all ridges of $SC(N,d)$ with a list of $2(d-1)$ consecutive vertices and every list of $2(d-1)$ consecutive vertices with a set of $(d-1)$ many ridges.

It follow that if we pick a tuple of $2(d-1)$ vertices $[v_1, ..., v_{2(d-1)}]$ uniformly at random and then pick one of the $(d-1)$ many ridges obtained by deleting a vertex from $[v_{d-1}, ..., v_{2(d-1)}]$ different from $v_{2(d-1)}$ this procedure generates a ridge uniformly at random from all ridges $\sigma_{i,j}$ with $i \geq d-2$ of $SC(N,d)$. For any fixed pattern π for a ridge we therefore have that under f the proportion of ridges $\sigma_{i,j}$ with $i \geq d-2$ colored by π is at most

$$(d-1)(d-1)!\frac{(c_1-(d-1))!}{(c_1-2(d-1))!}\frac{(1+\varepsilon)(c_1-2(d-1))!}{c_1!}\frac{1}{d-1}$$

This can be seen because we may specify first which ridge we will choose from a selected list of consecutive vertices, that is which vertex from among $v_{d-1}, ..., v_{2(d-1)-1}$ we will remove; there are $(d-1)$ choices. From here there will be $(d-1)!$ ways to map π to the selected positions and then $(c_1-(d-1))!/(c_1-2(d-1))!$ choices for the remaining positions to build a valid colored tuple of consecutive vertices. Each of these colored tuples will be selected with probability at most $(1+\varepsilon)(c_1-2(d-1))!/c_1!$ by the assumptions on f. But every tuple is associated with $d-1$ many ridges so we divide by $d-1$ to correct for this overcount.

Now we further simplify;

$$(d-1)(d-1)!\frac{(c_1-(d-1))!}{(c_1-2(d-1))!}\frac{(1+\varepsilon)(c_1-2(d-1))!}{c_1!}\frac{1}{d-1} = \frac{(1+\varepsilon)}{\binom{c_1}{d-1}}.$$

So no pattern on ridges will appear with proportion greater than $(1+\varepsilon)\Big/\binom{c_1}{d-1}$ among ridges $\sigma_{i,j}$ with $i \geq d-2$. Of course, this ignores patterns on ridges $\sigma_{i,j}$ with $i < d-2$ and $\{1, ..., d-1\}$, but this is a negligible proportion of ridges as N tends to infinity. The claim follows as we observe that $SC(N,d)$ has $1+(N-d+1)(d-1)$ ridges. $\square$

We can generalize this result for faces of any other codimension. This will be relevant for the proof of Theorem 3.1.3.

Corollary 3.2.5. *Fix d, and set $k \in \{1, \ldots, d-1\}$, $c_1 > 6(d-1)$, and $\varepsilon > 0$. There is a coloring of $SC(N,d)$ with c_1 colors such that all pattern classes of codimension k faces have at most size $(1+\varepsilon)N\binom{d-1}{k}\Big/\binom{c_1}{d-k}$ and no pair of intersecting $(d-k-1)$-faces have the same pattern for N large enough.*

Proof. We replicate the strategy of the proof of Proposition 3.2.3. As before, a randomized greedy algorithm picks a color that has not been used in the last $2(d-1)$ vertices as before, which is enough to guarantee that no pair of intersecting $(d-k-1)$-faces have the same pattern.

The Markov chain remains unchanged from the proof of Proposition 3.2.3; it's still defined on tuples of $2(d-1)$ colors. So each state corresponds to a coloring of $2(d-1)$ consecutive vertices, but now each state of S is associated to $\binom{d-1}{k}$ many $(d-k-1)$-faces that end at the final vertex of the consecutive list of $2(d-1)$ vertices. This comes from the fact that a $(d-k-1)$-face is a sequence of d consecutive vertices with k of them removed, but we do not remove the last one in order to associate each $(d-k-1)$-face to a unique state.

By the same reasoning as the proof of Proposition 3.2.3, which is the $k = 0$ case, we have that for N large enough, every pattern class will have at most $(1+\varepsilon)N\binom{d-1}{k} \Big/ \binom{c_1}{d-k}$ faces of codimension k. $\square$

3.2.2 The second coloring

We are now ready to refine a coloring of $SC(N, d)$ satisfying the conclusions of Proposition 3.2.3 so that no pair of ridges receive the same pattern. We do this via Proposition 3.2.6 which we state in a fairly general way to apply it to the pseudomanifold case later.

Proposition 3.2.6. *If C is a $(d-1)$-simplicial complex and there is a coloring of C on at most c_1 colors such that no pattern class of ridges has size more than S, no intersecting ridges receive the same pattern, and for any ridge σ, there are at most t other ridges which intersect σ then there is a refinement of the coloring with at most $c_1 \left\lceil \sqrt[d-1]{e(2tS+1)} \right\rceil$ colors and every ridge colored uniquely.*

As is the strategy in [59, 60], we will prove this proposition from the Lovász Local Lemma, originally due to Erdős and Lovász [23], which we recall in the symmetric version from [6] below:

Theorem 3.2.7 (Lovász Local Lemma [6])**.** *Let $A_1, A_2, \ldots, A_n$ be events in an arbitrary probability space. Suppose that each event A_i is independent of all other events A_j except for some subset of $\{A_1, \ldots, A_n\}$ with at most m elements. Suppose also that $\Pr[A_i] \leq p$ for all $1 \leq i \leq n$. If $ep(m+1) \leq 1$ then $\Pr[\bigwedge_{i=1}^n \overline{A_i}] > 0$*

Proof of Proposition 3.2.6. Let C be the complex described and colored by $f : V(C) \to [c_1]$ so that every pattern class of ridges has size at most S and no intersecting ridges receive the same pattern by f. We will find a second coloring g so that the final coloring satisfying the conclusions of the statements will be (f, g). The coloring g will be constructed randomly by choosing for each vertex a random color uniformly from a set of c_2 colors to be determined later.

For σ, τ, two ridges of C, with $f(\sigma) = f(\tau)$, let $A_{\sigma,\tau}$ denote the event that σ and τ receive the same color by (f, g). Clearly, if

$$\Pr\left(\bigwedge_{\substack{(\sigma,\tau) \text{ ridges of } C \\ \text{s.t. } f(\sigma)=f(\tau)}} \overline{A_{\sigma,\tau}} \right) > 0,$$

then there exists a choice for g on c_2 colors so that (f, g) colors every ridge uniquely. Indeed, for σ, τ with $f(\sigma) \neq f(\tau)$, refining the coloring by g will not cause σ and τ to receive the same pattern under (f, g). We apply the Local Lemma to the $A_{\sigma,\tau}$.

First, for disjoint σ, τ so that $f(\sigma) = f(\tau)$, there is a bijection $\phi : \sigma \to \tau$ sending each vertex in σ to the unique vertex in τ that receives the same coloring by f. Hence (f, g) gives σ and τ the same pattern if and only if $g(v) = g(\phi(v))$ for all $v \in \sigma$. Then we have the following bound on the probability of $A_{\sigma,\tau}$

$$\Pr(A_{\sigma,\tau}) \leq \left(\frac{1}{c_2}\right)^{d-1}.$$

We now bound for fixed (σ,τ), still with $f(\sigma) = f(\tau)$, the number of pairs $A_{\sigma',\tau'}$ so that $A_{\sigma',\tau'}$ is not independent of $A_{\sigma,\tau}$ *and so that* $f(\sigma') = f(\tau')$ also. Clearly if $(\sigma \cup \tau) \cap (\sigma' \cup \tau') = \emptyset$, then $A_{\sigma,\tau}$ is independent of $A_{\sigma',\tau'}$. For (σ,τ) fixed we bound from above the number of pairs (σ',τ') so that $(\sigma \cup \tau) \cap (\sigma' \cup \tau') \neq \emptyset$. If $(\sigma' \cup \tau')$ is to intersect $(\sigma \cup \tau)$ then without loss of generality we may assume that σ' intersects either σ or τ. There are at most t choices for σ' so that σ' intersects σ, and from here τ' may be selected to be any ridge with the same pattern as σ', and there are S such ridges. So for (σ,τ) there are at most tS choices for (σ',τ') so that σ' intersects σ and another at most tS choices so that σ' intersects τ. Therefore, $A_{\sigma,\tau}$ is independent of all but at most $2tS$ other events $A_{\sigma',\tau'}$.

It follows that choosing c_2 large enough that

$$e\frac{1}{c_2^{d-1}}(2tS+1) \leq 1,$$

will imply by the Lovász Local Lemma that with positive probability $A_{\sigma,\tau}$ fails to hold simultaneously for all pairs of ridges σ,τ with $f(\sigma) = f(\tau)$. By setting $c_2 = \left\lceil \sqrt[d-1]{e(2tS+1)} \right\rceil$, we complete the proof. □

With Proposition 3.2.3 and 3.2.6 now proved we are ready to prove Theorem 3.1.2.

Proof of the lower bound for Theorem 3.1.2. Let $d \geq 3$ be fixed. The result is asymptotic in n, so here we fix $c_1 > 6(d-1)$ and $\varepsilon > 0$, and we will show the constant factor in the lower bound emerges as $c_1 \to \infty$ and $\varepsilon \to 0$.

Let N be large enough so that there exists a coloring of $SC(N,d)$ with c_1 colors so that all pattern classes of ridges have size at most $S = (1+\varepsilon)(d-1)N \Big/ \binom{c_1}{d-1}$, and let $f : V(SC(N,d)) \to [c_1]$ be such a coloring. To apply Proposition 3.2.6 we need to find an upper bound for the maximum number of ridges that intersect any given ridge, that is in the notation of the proposition we should find a suitable value for t in the case of $SC(N,d)$.

The facets of $SC(N,d)$ are d consecutive vertices in $[N]$, so the ridges are obtained by taking any d consecutive elements of $[N]$ and removing exactly one element. For any fixed ridge σ, there is a facet $F := F(\sigma)$ so that σ is contained in F. From the structure of $SC(N,d)$, we have that F intersects at most $2d$ other facets and each of these facets contains d ridges. Thus σ intersects at most $2d^2$ other ridges.

With c_1, S, and t determined we apply Proposition 3.2.6 to say that there is a coloring of $SC(N,d)$ by at most c_1c_2 colors where $c_2 = \left\lceil \sqrt[d-1]{e(2tS+1)} \right\rceil$ so that every ridge has a unique pattern. Let $(f,g) : V(SC(N,d)) \to [c_1c_2]$ be such a coloring with $t = 2d^2$. By Lemma 3.2.2 the complex $C := SC(N,d)/(f,g)$ has diameter equal to the diameter of

$SC(N,d)$. Therefore, the diameter of C is $N-d$, and the number of vertices of C is at most

$$\begin{aligned} c_1c_2 &= c_1\left\lceil \sqrt[d-1]{e(2tS+1)}\right\rceil \\ &= c_1\left\lceil \sqrt[d-1]{e\left(2(2d^2)\frac{(1+\varepsilon)(d-1)N}{\binom{c_1}{d-1}}+1\right)}\right\rceil. \end{aligned}$$

Now as c_1 tends to infinity $c_1^{d-1}\Big/\binom{c_1}{d-1}$ tends to $(d-1)!$. Thus, given $\delta>0$ we may set ε small enough and c_1 large enough so that for all N large enough we have a complex on at most $(1+\delta)\sqrt[d-1]{(d-1)!4d^3(N-d)e}$ vertices with diameter $N-d$. Letting n be the number of vertices in this complex we have a complex on n vertices with diameter at least

$$\left(\frac{1}{1+\delta}\right)^{d-1}\frac{n^{d-1}}{4ed^2d!}.$$

This proves the theorem as δ is arbitrary. □

3.3 Pseudomanifold case

Here we prove Theorem 3.1.3. In this case we establish a tighter upper bound than in the general case based on the observation that the dual graph of a $(d-1)$-dimensional pseudomanifold is a d-regular graph.

For arbitrary regular graphs there is the following bound on the diameter; this result is a special case of Theorem 5 of Caccetta and Smyth [12] and also appears in a different form earlier in work of Moon [58].

Theorem 3.3.1 ([58, 12])**.** *Let G be a connected d-regular graph on n vertices. Then*

$$\operatorname{diam}(G)\le\frac{3n}{d+1}.$$

From this it follows that we can give an upper bound to the combinatorial diameter of a pseudomanifold as in Theorem 3.1.3

Proof of the upper bound for Theorem 3.1.3. Here we use the notation f_i for a complex to denote the usual f-vector entry. Let C be a $(d-1)$-dimensional pseudomanifold on n vertices, then $\mathrm{G}(C)$ is a d-regular graph on $f_{d-1}(C)$ vertices. By Theorem 3.3.1 we have that

$$\operatorname{diam}(\mathrm{G}(C))\le\frac{3f_{d-1}(C)}{d+1}.$$

Now, since C is a $(d-1)$-dimensional pseudomanifold on n vertices,

$$df_{d-1}(C)=2f_{d-2}(C)\le 2\binom{n}{d-1}.$$

Therefore,

$$\operatorname{diam}(\mathrm{G}(C)) \leq \frac{6}{d(d+1)}\binom{n}{d-1} \leq \frac{6n^{d-1}}{(d+1)!}.$$

□

We prove the lower bound of Theorem 3.1.3 by the same method as Theorem 3.1.2 with a different starting complex. For given d, we will start with the top-dimensional boundary of $SC(N, d+1)$, denoted $\partial SC(N, d+1)$ and color the vertices as in Lemma 3.2.2. It is clear that $\partial SC(N, d+1)$ is a pseudomanifold. In fact it is a triangulated sphere as it has the combinatorial type of a stacked polytope.

We now proceed to describe the facets of $\partial SC(N, d+1)$. These facets are the ridges of $SC(N, d+1)$ that are contained in exactly one facet of $SC(N, d+1)$. Recall that facets of $SC(N, d+1)$ are sequences of $d+1$ consecutive vertices. Then, ridges of $SC(N, d+1)$ are sequences of $d+1$ vertices with one of them missing. If we remove the first or the last vertex of a facet, then the corresponding ridge of $SC(N, d+1)$ is contained in two facets, unless it is the first d vertices or the last d vertices. This gives us three types of facets of $\partial SC(N, d+1)$. The first two types each correspond to a single facets and are:

$$\alpha = \{1, ..., d\},$$

and

$$\omega = \{N-d+1, ..., N\}.$$

We call these facets of $\partial SC(N, d+1)$ respectively the *initial facet* and the *final facet*. We refer to the remaining facets of $\partial SC(N, d+1)$ as *middle facets*. These are given by

$$\tau_{i,j} = \{i, i+1, i+2, \ldots, i+d\} \setminus \{i+j\},$$

for $i \in [N-d]$ and $j \in \{1, \ldots d-1\}$,

From now on we are working with $\partial SC(N, d+1)$, so when we talk about *facets* or *ridges*, we are referring to that complex; facets are of dimension $d-1$ and ridges are of dimension $d-2$.

The diameter of $\partial SC(N, d+1)$ is less obvious than that of $SC(N, d)$, so we first prove the following lemma to establish a lower bound on the diameter of $\partial SC(N, d+1)$.

Lemma 3.3.2. *The diameter of $\partial SC(N, d+1)$ is at least $\left(\frac{d-1}{d}N\right) - d$.*

Proof. We will use a potential function to provide a lower bound on the length in the dual graph of any path from α to ω. Observe that any path from α to ω will only have middle facets between α and ω. We set a potential function p over the middle facets of $\partial SC(N, d+1)$, defined by $p(\tau_{i,j}) = i - \frac{j}{d-1}$. We will show that every move from one middle facet to an adjacent middle facet increases this potential at most by $1 + \frac{1}{d-1}$.

A move in the dual graph corresponds to removing one vertex from a facet, and adding a new vertex. The set of vertices shared by the initial and final facets then form the ridge in $\partial SC(N, d+1)$ connecting the two.

Since $\partial SC(N, d+1)$ is a pseudomanifold, every facet is adjacent to exactly d facets, corresponding to d choices of a vertex to remove. Moreover the new vertex to add is uniquely determined because in a pseudomanifold every ridge is in exactly two facets. For a move between middle facets along a path from α to ω there are three cases to consider:

- We remove the last vertex of $\tau_{i,j}$. In this case, the adjacent facet is either $\tau_{i-1,j+1}$ if $j \neq d-1$, or $\tau_{i-2,1}$ if $j = d-1$. In both cases, the potential decreased.
- We remove the first vertex. If $j = 1$, the adjacent facet is $\tau_{i+2,d-1}$. If $j \neq 1$, the next facet is $\tau_{i+1,j-1}$. In both cases, the increase in potential is $1 + \frac{1}{d-1}$.
- We remove a vertex $i + j' \in \{i+1, \ldots, i+d-1\}$ different from the first or the last vertex. Then, the adjacent facet is $\tau_{i,j'}$, and the potential increased at most by $\frac{d-1}{d-1} = 1$.

From these three cases we see that any step between middle facets of $\partial SC(N, d+1)$ increases the potential at most by $1 + \frac{1}{d-1} = \frac{d}{d-1}$. Observe that α is adjacent only to facets with potential at most one, since its neighborhood is $N(\alpha) = \{\tau_{1,j} : j = 1 \ldots d-1\} \cup \{\tau_{2,d-1}\}$. Similarly, the neighborhood of ω is $N(\omega) = \{\tau_{N-d,j} : j = 2, \ldots, d\} \cup \{\tau_{N-d-1,1}\}$, so ω is adjacent only to facets with potential at least $N - d - 2$.

Any path from α to ω passes only through middle facets, and increases the potential from a starting value on the first middle facet at most 1 to a final value on the final middle facet of at least $N - d - 2$ potential. And every step between these facets increases the potential at most by $\frac{d}{d-1}$. Hence, the middle part of the path takes at least $((N-d-2)-1)\,\frac{d-1}{d}$ steps. And the path from α to ω is at least two steps longer than that.

Therefore, the shortest paths between α and ω have length are at least $((N-d-2)-1)\frac{d-1}{d} + 2 = \frac{d-1}{d}N - d + \frac{3}{d} > (d-1)N/d - d$. The diameter of $\partial SC(N, d+1)$, as the largest distance between any pair of vertices, is at least as large. $\square$

This lower bound on the diameter of $\partial SC(N, d+1)$ is actually sharp for N large enough, but we do not need to prove that for the purposes of this article.

Now we want to color $\partial SC(N, d+1)$ as in the proof of the general case to apply Lemma 3.2.2.

Proof of the lower bound for Theorem 3.1.3. As in the proof of Theorem 3.1.2, fix $d \geq 2$ as well as $c_1 > 6(d-1)$ and $\varepsilon > 0$. We let N be large enough so that we may apply Corollary 3.2.5 to color $SC(N, d+1)$ by a coloring f with c_1 colors so that all pattern classes of $(d-2)$-dimensional faces have size at most

$$S = \frac{(1+\varepsilon)N\binom{(d+1)-1}{2}}{\binom{c_1}{(d+1)-2}}.$$

This coloring then induces a coloring of $\partial SC(N, d+1)$ so that no pattern class of ridges of $\partial SC(N, d+1)$ has size more than S and no intersecting ridges receive the same pattern.

In order to apply Proposition 3.2.6 to $\partial SC(N, d+1)$, we need to determine an upper bound t for the maximum number of ridges that intersect any given ridge. We will show now that $t = (d+1)^3$ is such a bound.

Recall that facets of $\partial SC(N, d+1)$ are sequences of $d+1$ consecutive vertices with one of them removed. Then ridges are sequences of $d+1$ consecutive vertices with two of them removed. For a fixed ridge, the number of choices of potential such intervals is at most $2d+1$. And for each choice of an interval of consecutive vertices, we have to remove two of them. This gives us a total of at most $(2d+1)\binom{d+1}{2} \leq (d+1)^3$ ridges intersecting the fixed one.

With c_1, S, and t determined we apply Proposition 3.2.6 to color $\partial SC(N, d+1)$. The number of colors required is at most

$$c_1 \left\lceil \sqrt[d-1]{e\left(2(d+1)^3 \frac{(1+\varepsilon)N\binom{d}{2}}{\binom{c_1}{d-1}} + 1\right)} \right\rceil.$$

The proposition guarantees that no pair of ridges receive the same pattern.

Thus, exactly as in the proof of Theorem 3.1.2, for any $\delta > 0$ we set c_1 large enough and ε small enough so that for N large enough have a complex on at most

$$(1+\delta)\sqrt[d-1]{e(d-1)!(d+1)^3 N d^2}$$

vertices, which has diameter at least

$$(1-\delta)\frac{d-1}{d}N.$$

Letting n denote the number of vertices in this complex we have a complex on n vertices whose diameter is at least

$$\begin{aligned} & (1-\delta)\left(\frac{1}{1+\delta}\right)^{d-1} \frac{(d-1)n^{d-1}}{e(d+1)^3 d^3 (d-1)!} \\ \geq \quad & (1-\delta)\left(\frac{1}{1+\delta}\right)^{d-1} \frac{n^{d-1}}{4ed^4 d!}. \end{aligned}$$

As δ is arbitrary, this proves the claim. $\square$

Chapter 4

Width of topological prismatoids

This chapter is from the article "Topological prismatoids and small simplicial spheres of large diameter" [17], by Francisco Criado and Francisco Santos.

This is an Accepted Manuscript of an article published by Taylor & Francis in Experimental Mathematics on 22 Jul 2019, available online: `https://www.tandfonline.com/doi/abs/10.1080/10586458.2019.1641766`.

4.1 Introduction

The main result of this chapter is the construction of non-Hirsch simplicial $(d-1)$-spheres with $d = n - d = 9$, that is, simplicial $(d-1)$-spheres with diameter higher than the Hirsch bound:

Theorem 4.1.1. *There exist simplicial* 8*-spheres with* 18 *vertices and diameter at least* 9*. That is,*

$$H_{sph}(18, 9) > 9.$$

These values for n and d are close to minimal since the inequality $H_{\mathrm{sph}}(n, d) \leq n - d$ is known to hold for $n - d \leq 5$. This was proved for polytopal spheres by Klee and Walkup [46], and we show in Theorem 4.2.1 how to modify their proof for non-polytopal ones.

One reason to concentrate on small examples is that from them it is very easy to construct higher dimensional ones by application of Lemma 1.2.1 In particular, from the example mentioned in Theorem 4.1.1 one easily derives that $H_{\mathrm{sph}}(2d, d) > d$ for every $d \geq 9$.

Adding connected sums to the suspensions used in this proof we obtain a more refined asymptotic lower bound (see Theorem 1.2.14). If we apply it to the non-Hirsch sphere of Theorem 4.1.1 we have that:

Corollary 4.1.2. *For every n and d,*

$$H_{sph}(n,d) > \left\lfloor \frac{n-d}{d} \right\rfloor \cdot \left(\left\lfloor \frac{10d}{9} \right\rfloor - 1 \right) \simeq 1.11(n-d).$$

Our construction leading to Theorem 4.1.1 uses the same *prismatoid* technique developed by Santos [63, 55] and present in all non-Hirsch polytopes known so far, but we abstract it to a combinatorial/topological context. Recall (from Section 1.2.3) that a d-prismatoid P is a d-dimensional polytope whose vertices lie in two parallel facets. Removing from ∂P the relative interiors of these two facets produces a polyhedral complex homeomorphic to the Cartesian product of a $(d-2)$-sphere and a segment, and with all its vertices in the boundary, that we assume simplicial. We call a simplicial complex with these topological properties a *topological $(d-1)$-prismatoid*. The *width* of a topological prismatoid $\mathcal{C}$ is defined to equal 2 plus the minimum distance, in the adjacency graph, between facets incident to one and the other component of $\partial\mathcal{C}$. In this setting we prove the following analogue of [63, Theorem 2.6]:

Theorem (Theorem 4.3.5)**.** *Let $\mathcal{C}$ be a topological prismatoid of dimension $(d-1)$, of width l and with $n > 2d$ vertices. Assume that its two bases are polytopal. Then, there exists a topological $(n-d-1)$-prismatoid $\mathcal{C}'$ with $2n-2d$ vertices and width at least $l+n-2d$.*

In particular, if $l > d$ then $\mathcal{C}'$ is a non-Hirsch simplicial sphere of dimension $D-1 := n-d-1$, with $N := 2D = 2n-2d$ vertices, and of adjacency diameter larger than $N-D$.

The condition on the politopality of the bases is purely technical, and we believe a weaker condition could work. But this one suffices for our purposes later on.

Since this theorem is related to the d-step property of Klee and Walkup [46], we say that topological $(d-1)$-prismatoids of width larger than d are *non-d-step*. Our goal is to find non-d-step topological prismatoids with $n-d$ as small as possible.

We do this by a *simulated annealing* approach on the graph of *flips* among non-d-step prismatoids of a given dimension. That is, we start with a topological $(d-1)$-prismatoid of width $l > d$ and do flips in it at random, but preserving the width and giving higher probability to the flips that go in the direction of decreasing n. See Section 4.3.1 for the definition and properties of flips in topological prismatoids, and Section 4.4 for details of our heuristics and implementation.

We choose as starting point the 28-vertex prismatoid constructed in [55, Corollary 2.9], which has dimension $d = 5$ as a polytope, that is, $d-1 = 4$ as a topological prismatoid. From it we find thousands of non-d-step topological 4-prismatoids with number of vertices below 28. In particular, we find four combinatorially different ones with 14 vertices. Any of these four implies Theorem 4.1.1, via Theorem 4.3.5. Observe that although the prismatoids are found computationally, the proof that they are non-d-step is elementary and can be done by hand.

The constructed prismatoids are analyzed a bit in Section 4.5. For the four smallest ones we have checked that they are *shellable*, with respect to a natural notion of shelling of topological prismatoids that we introduce in Section 4.3.4, and which implies shellability for the resulting non-Hirsch spheres mentioned in Theorem 4.1.1 (see Proposition 4.3.14). Shellability is necessary for polytopality, but we do not know whether our prismatoids (or spheres) are polytopal.

4.2 A lower bound to the complexity of non-Hirsch simplicial spheres

In this section we adapt a result for the lower bound of $H(n, d)$ when $n - d \leq 5$ to the topological context. This shows that our sphere, with $n - d = 9$, can not be far from the optimum.

Theorem 4.2.1 (Sphere version of [46, Theorem 4.2])**.** *For any n, d with $n - d \leq 5$, $H_{sph}(n, d) \leq n - d$.*

Proof. By Proposition 1.2.11 we can assume $n = 2d$. So, let $\mathcal{S}$ be a $(d-1)$-sphere with $2d$ vertices and let F and G be two facets in it achieving the diameter. We want to show that he distance from F to G to be at most d.

The proof follows the one in [46] for the polytopal case and consists of the following three claims. Claims 1 and 2 do not need the assumption that $d \leq 5$, and correspond to Theorem 2.8 and Proposition 3.4.(b) in [46].

Claim 1: there is no loss of generality in assuming $F \cap G = \emptyset$. If this is not the case, let $\mathcal{S}' = \mathrm{link}_{\mathcal{S}}(F \cap G)$. Then, $F' := F \setminus G$ and $G' := G \setminus F$ are facets in $\mathcal{S}'$ at distance at least equal to the distance between F and G in $\mathcal{S}$. Observe that $\mathcal{S}'$ is a $(d-1-k)$-sphere, where $k = |F \cap G|$, and it has between $2d - 2k$ and $2d - 1$ vertices. If $\mathcal{S}'$ has $2d - 2k$ vertices, the diameter of $\mathcal{S}'$ is bounded by $H_{\mathrm{sph}}(2d - 2k, d - k)$ which, by induction on d, is at most $d - k$. If $\mathcal{S}'$ has more than $2d - 2k$ vertices, let $\mathcal{S}''$ be the iterated one-point suspension of $\mathcal{S}'$ at each vertex w not in $F' \cup G'$. $\mathcal{S}''$ is a sphere of dimension $d - k - 1 + l$ and with $2d - 2k + 2l$ vertices, where l is the number of times we did the one-point suspension. In $\mathcal{S}'$, F' and G', each together with one of the two vertices of each suspension, give two complementary facets which are at least at the same distance as F' and G' were in $\mathcal{S}'$.

Claim 2: assume $F \cap G = \emptyset$. There are vertices $v \in F$ and $w \in G$ such that $\{v, w\}$ is an edge in $\mathcal{S}$ and such that F and G are adjacent respectively to facets F' and G' with $\{u, v\} \in F' \cap G'$. Let F' be any facet adjacent to F. Let w be the unique vertex in $F' \setminus F$ and u the unique vertex in $F \setminus F'$. It is impossible for the $d - 1$ facets adjacent to G and containing w to all of them use u: indeed, if that happened then these $d - 1$ facets together with G form a ball with w in its interior, implying that there are no other facets in the star of w; this is impossible because F' is in the star of w too. Hence, there is a facet G' adjacent to G, using w, and using a vertex v of F other than u. This proves the claim for the edge vw thus obtained.

Claim 3: with v, w as above, there is a path in $\mathrm{star}_{\mathcal{S}}(\{v, w\})$ of length at most $d - 2$ between a facet F' adjacent to F and a facet G' adjacent to G. The proof of this claim is the complicated part, occupying most of Section 4 (pages 69–71) of [46]. The good thing is that we do not to even check that the proof extends to spheres, since it is a statement about the link of the edge $\{u, v\}$ in $\mathcal{S}$ and that link is a $(d-3)$-sphere, hence polytopal for $d \leq 5$. Indeed, what Klee and Walkup show is that the following (which paraphrases [46, p. 69, ll. 4–8]) holds for every simplicial k-polytope Q with $k \leq 3$ and with between $2k$ and $2k+2$ vertices: "if the vertices of Q are divided into two disjoint classes X and Y, each consisting of at most $k + 1$ vertices, then there is a path of length at most k from a facet contained in X and a facet contained in Y." This proves the claim by letting Q be a polytope isomorphic to $\mathrm{link}_{\mathcal{S}}(\{v, w\})$, $k = d - 2$, and X and Y be $F \setminus \{v\}$ and $G \setminus \{w\}$, respectively. $\square$

4.3 Topological prismatoids and the topological strong d-step theorem

4.3.1 Prismatoids and flips in them

We now define the main object we work with:

Definition 4.3.1. *A* (($d-1$)-dimensional) topological prismatoid $\mathcal{C}$ *is a* $(d-1)$*-dimensional pure simplicial complex homeomorphic to* $\mathbb{S}_{d-2} \times [0,1]$ *(that is, it is homeomorphic to a cylinder), and such that every face with all its vertices in the same boundary component is a boundary face. Put differently, the two boundary components, each homeomorphic to* $\mathbb{S}_{d-2}$*, are induced subcomplexes.*

Bistellar flips were introduced for general manifolds in [62] and they are a standard tool in combinatorial topology by now, as local modifications that preserve the PL-type. The main result of Pachner [62] is the converse: every two PL-homeomorphic simplicial manifolds can be transformed into one another via a sequence of bistellar flips. We here adapt the general definition to the case of topological prismatoids:

Definition 4.3.2. *A* flip *in a topological prismatoid* $\mathcal{C}$ *is a triple* (f, l, v) *of pairwise disjoint subsets of* $V(\mathcal{C})$ *such that* f *is a face,* l *is a minimal non-face,* $\mathrm{link}_{\mathcal{C}}(f) = \partial(l) * v$*,* v *is either the empty face or a vertex, and one of the following two things happens:*

- $|f| + |l| = d + 1$ *and* $v = \emptyset$*, in which case* l *is required to intersect both bases of* $\mathcal{C}$*. (Observe that in this case* $\mathrm{link}_{\mathcal{C}}(f) = \partial(l) * v = \partial(l)$*).*
- $|f| + |l| = d$ *and* v *is a vertex, in which case* f *and* l *are required to be contained in the base opposite to* v*.*

In both cases, the result of the flip is the prismatoid

$$\mathcal{C}' = \mathcal{C} \setminus \mathrm{star}_{\mathcal{C}}(f) \cup (l * \partial(f) * v).$$

flips with $v = \emptyset$ *are called* interior flips *and flips where* v *is a vertex are called* boundary flips*. The* support *of the flip is* $f \cup l \cup v$*.*

Put differently, an (f, l, v) flip removes all faces containing f and inserts as new faces all subsets of $f \cup l \cup v$ that contain l but not f. The interior flips do not change the boundary and are exactly the traditional bistellar flips; the main new feature of our definition is that we allow flips that change the boundary, but guaranteeing the two bases to still be induced subcomplexes after the flip. Observe that a boundary flip can remove or add a vertex. This happens when f or l, respectively, have size 1.

Remark 4.3.3. *In Definition 4.3.2,* $V(\mathcal{C})$ *is understood as the* ground set *of the prismatoid, which may contain points that are not used as vertices. In particular, a boundary flip may have* $l = \{w\}$ *for a* w *that is not a vertex, and* f *a facet in a base. The result of the flip is that this facet is stellarly subdivided with the new vertex* w*.*

We need this type of flips because we want flips to be reversible for the simulated annealing framework. These flips are the inverse of vertex-removing flips.

An important feature used in our implementation is that knowing only the support u of a flip we can recover the sets f, l and v and thus perform the flip:

- If u has a single vertex from one of the bases then the flip is a boundary flip, and that vertex is v. Indeed, in an interior flip l has at least one vertex from each component by definition, and f has at least another from each base because the condition $\mathrm{link}(f) = \partial(l)$, with $|f| + |l| = d + 1$, implies that f is an interior face.
- In both cases, the set $f \cup v$ equals the intersection of all facets of $\mathcal{C}$ contained in u. This allows us to recover f, and hence l, once we know v by the previous point.

The support u of a flip must have $d + 1$ vertices, since it is the vertex set of a d-ball of the form $l * \partial(f) * v$, that is, the join of an i-simplex and the boundary of a j-simplex, with $i + j = d - 1$.

Moreover, the following result allows us to detect flips:

Proposition 4.3.4. *Given a set u of $d+1$ vertices (or d vertices and an unused point in the case of insertion flips) not all in one base, let f, l and v be as above. The following conditions are necessary and sufficient for u to support a flip in (f, l, v):*

1. *u is the neighborhood of a ridge.*
2. *neigh(f) has $d + 1$ vertices (or d vertices in case of insertion flips).*
3. *l is not a face of $\mathcal{C}$.*

4.3.2 The strong d-step theorem for topological prismatoids

We here prove the main theoretical result that allows us to use topological prismatoids to search for non-Hirsch spheres. The *width* of $\mathcal{C}$ is two plus the distance, in the adjacency graph, between the set of facets incident to one base and the set of facets incident to the other (the distance between two sets is, as customary, the minimum distance between respective elements).

Theorem 4.3.5 (Strong d-step theorem for topological prismatoids)**.** *Let $\mathcal{C}$ be a topological prismatoid of dimension $(d - 1)$, width l and with $n > 2d$ vertices. Assume that its two bases are polytopal. Then, there exists a topological $(n - d - 1)$-prismatoid $\mathcal{C}'$ with $2n - 2d$ vertices and width at least $l + n - 2d$.*

In particular, if $l > d$ then $\mathcal{C}'$ is a simplicial sphere of dimension $D - 1 := n - d - 1$, with $N := 2D = 2n - 2d$ vertices whose adjacency graph has diameter larger than $N - D$.

Remark 4.3.6. The excess of the non-Hirsch sphere produced via Theorem 4.3.5 from a topological $(d - 1)$-prismatoid $\mathcal{C}$ of width l and n vertices equals

$$\frac{l - d}{n - d}.$$

Thus, we call that quotient the *(prismatoid) excess* of $\mathcal{C}$.

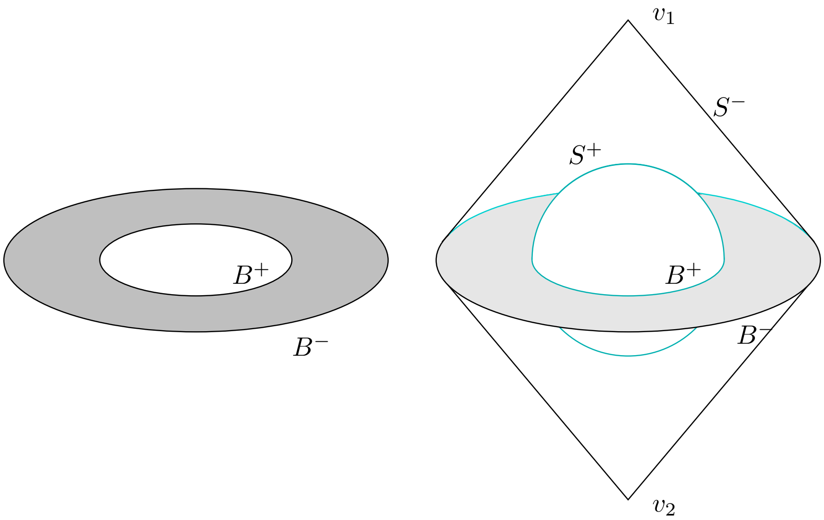

Figure 4.1: Sketch of the construction of the prismatoid $\mathcal{C}'$ (right) from $\mathcal{C}$ (left) in the proof of Theorem 4.3.5

Proof. By induction on $n - 2d$ it suffices to construct a prismatoid of dimension d with $n+1$ vertices, width at least $l+1$, and such that its bases are polytopal. Repeating this procedure $n - 2d$ times we arrive at a $(D-1)$-prismatoid with $2D$ vertices. In such a prismatoid the bases are simplices, so the prismatoid is a $(D-1)$-sphere with $2D$-vertices and diameter at least $l + (n - 2d)$.

For the inductive step, let B^+ and B^- be the two bases of $\mathcal{C}$. Since $\mathcal{C}$ has more than $2d$ vertices, at least one of them (say B^+) is not a simplex. Let S^+ be a simplicial polytopal $(d-1)$-sphere containing B^+ and with no additional vertices. S^+ exists since B^+ is polytopal: Let $P \subset \mathbb{R}^d$ be a d-polytope realizing B^+ and choose a sufficiently generic lifting function $h : \text{vertices}(P) \to \mathbb{R}$. Then S^+ can be chosen to be the boundary complex of $\text{conv}\{(v, h(v)) : v \in \text{vertices}(P)\}$.

Let S_1^+ and S_2^+ be the two closed $(d-1)$-balls whose intersection is B^+ and whose union is S^+. Let v_1 and v_2 be two additional vertices. Consider the following simplicial complex:

$$\mathcal{C}' := (\mathcal{C} \cup B_1^+) * v_1 \cup (\mathcal{C} \cup B_2^+) * v_2.$$

$\mathcal{C}'$ is a topological d-prismatoid with bases S^+ and the suspension $S^- := B^- * \{v_1, v_2\}$ of B^- (see Figure 4.1). It is not yet the prismatoid we want since it has $n+2$ vertices and we want only $n+1$. Later in the proof we show how to reduce the number of vertices by one, but for the time being let us not care about that. Instead, for reasons that will become apparent later, when computing the length of paths in the adjacency graph of $\mathcal{C}'$ we will neglect the steps of the form $F * v_1$ to $F * v_2$ or vice versa, for facets F of $\mathcal{C}$. We claim that, even with this reduced way of counting steps, $\mathcal{C}'$ has width strictly larger than $\mathcal{C}$.

To prove the claim, let $F_0', \ldots, F_t'$ be a path in $\mathcal{C}'$ from a facet F_0' adjacent to S^- to a facet F_t' adjacent to S^+. Since we want our path as short as possible, there is no loss of generality in assuming that F_t' is the only facet adjacent to S^+. That is, each facet

F'_i, $i \in \{0, \dots, t-1\}$, is of the form $F_i * v_j$ for a certain facet F_i of $\mathcal{C}$ and $j \in \{1, 2\}$. We have that:

- Facets F_i and F_{i+1} are either adjacent or the same; the latter happens if and only if F'_i and F'_{i+1} are of the form $F * v_1$ and $f * v_2$ for the same facet F of $\mathcal{C}$.
- The first facet F_0 is adjacent to B^-, since F'_0 is adjacent to $S^- = B^- * \{v_1, v_2\}$.
- The last facet F_{t-1} is adjacent to B^+, since F'_{t-1} is obtained from the facet $F'_t = F_t * v_j$ by changing a single vertex, and F_t is a facet of S^+, not of $\mathcal{C}$.

Thus, as claimed, the width of $\mathcal{C}$ is strictly larger than that of $\mathcal{C}$, even neglecting the steps $F * v_1 \leftrightarrow F * v_2$.

We now get rid of one vertex without decreasing the width of $\mathcal{C}'$. For this, let v be any vertex of B^- and observe that

$$\mathrm{link}_{\mathcal{C}'}(vv_1) = \mathrm{link}_{\mathcal{C}'}(vv_2) = \mathrm{link}_{\mathcal{C}}(v).$$

This implies that we can substitute in $\mathcal{C}'$ the stars of edges vv_1 and vv_2 by the star of a single edge v_1v_2, to obtain a topological prismatoid with one less vertex, that is, with $n+1$ vertices. More precisely, we let:

$$\mathcal{C}'' := \mathcal{C}' \setminus (vv_1 * \mathrm{link}_{\mathcal{C}}(v) \cup vv_2 * \mathrm{link}_{\mathcal{C}}(v)) \cup v_1v_2 * \mathrm{link}_{\mathcal{C}}(v).$$

restricted to S^-, which was the suspension of B^-, this operation produces the one point suspension of B^- at vertex v. Thus, $\mathcal{C}''$ is a prismatoid with bases S^+ and the one-point suspension of B^- at v. It has $n+1$ vertices, dimension d, and it has the same width as $\mathcal{C}'$ if we neglect the steps of the form $F*v_1 \leftrightarrow F*v_2$; in particular, it has width strictly larger than that of $\mathcal{C}$. S^+ is a polytopal $(d-1)$-sphere by the way we constructed it, and the other base is also polytopal since one-point suspensions of polytopes are polytopes. ☐

Remark 4.3.7. Analyzing the proof of Theorem 4.3.5 the reader can check that the only place where we need the bases of $\mathcal{C}$ to be polytopal is in order to construct the $(d-1)$-sphere S^+ from the base B^+ of $\mathcal{C}$. That is, strictly speaking we do not need each base B to be polytopal but only to be embeddable in a sphere of one more dimension without extra vertices (and we need to keep this property recursively until B becomes a simplex). We do not know whether all spheres have this property but we suspect not.

4.3.3 Prismatoids of large width via reduced incidence patterns

In all previous constructions of non-Hirsch polytopes, the proof that the prismatoids to which Theorem 4.3.5 is applied is non-d-step uses the following result.

Proposition 4.3.8. *Let $\mathcal{C}$ be a (geometric or topological) prismatoid with bases B^+ and B^-. A necessary condition for $\mathcal{C}$ to be d-step is that there are vertices $v \in B^+$ and $w \in B^-$ such that vw is an edge and the star of vw contains facets incident to both bases.*

This proposition is a rephrasing of [55, Proposition 2.1], and used in part (3) of [63, Lemma 5.9]. Observe that the necessary condition in the statement is exactly Claim 2 in the proof of Theorem 4.2.1. Following [55] we introduce the following graph-theoretical way to visualize this property:

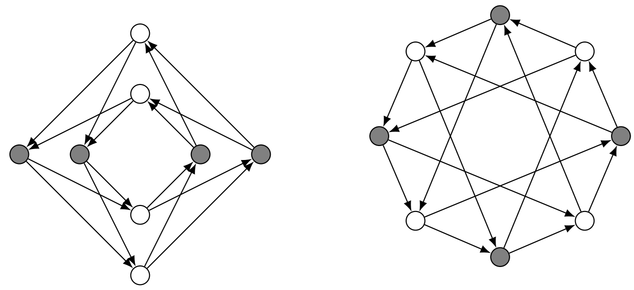

Figure 4.2: The two minimal reduced incidence patterns without cycles of length two

Definition 4.3.9. *Let $\mathcal{C}$ be a topological prismatoid with bases B^+ and B^-. The* incidence pattern *of $\mathcal{C}$ is the bipartite directed graph having a node for each vertex of $\mathcal{C}$ with bipartition given by the bases and with the following arcs: for each $v \in B^+$ and $w \in B^-$ we have an arc $v \to w$ (resp., $w \to v$) if there is a facet F in $\mathcal{C}$ containing vw and incident to B^+ (resp., incident to B^-). The* reduced incidence pattern *is the subgraph induced by vertices that are not sources.*

In this language, Proposition 4.3.8 becomes:

Proposition 4.3.10 (Topological version of [55, Proposition 2.3])**.** *Let $\mathcal{C}$ be a topological prismatoid. If there is no directed cycle of length two (that is, a "bidirectional arc") in its reduced incidence pattern then $\mathcal{C}$ is non-d-step.*

Proof. Suppose $\mathcal{C}$ is d-step, so that there is a sequence of facets $F_1, \ldots, F_{d-1}$, each adjacent to the next, and with F_0 adjacent to B^+ and F_1 adjacent to B^-. In particular, F_1 consists of a vertex v of B^- and $d-1$ vertices of B^+, and F_{d-1} consists of a vertex w of B^+ and $d-1$ vertices of B^-. This can only happen if v is already in F_{d-1} and w in F_1, so that v and w form a 2-cycle in the reduced incidence pattern. $\square$

Remark 4.3.11. *The absence of cycles of length two is sufficient but not necessary for being non-d-step. For example, two of the four small non-Hirsch prismatoids described in Section 4.5 do have cycles of length two in their reduced incidence patterns.*

Using the fact that in a reduced incident pattern without cycles of length two all vertices must have out-degree at least two, the minimum possible patters were classified in [55]. The proof works without changes for topological prismatoids:

Lemma 4.3.12 ([55, Proposition 2.4])**.** *Let $\mathcal{C}$ be a topological prismatoid whose reduced incidence pattern has no cycles of length two. Then, the reduced incident pattern has at least eight vertices.*

Moreover, the only two possible patterns with eight vertices are those of Figure 4.2 (vertices of one base are white, vertices of the other are grey).

It is interesting to note that the prismatoids constructed in [55, 63] have the reduced incidence pattern on the left of Figure 4.2 while the ones we obtain in this paper are related to the pattern on the right. See Section 4.5 for details.

4.3.4 Shellability of (topological) prismatoids

The following concept of shelling for prismatoids is a special case of shelling of *relative simplicial complexes*, in the sense of [32, Section 4.2]. Indeed, $\mathcal{C}$ is shellable from B^+ to B^- in the sense of Definition 4.3.13 if and only if the relative complex $(\mathcal{C}, B^+)$ is shellable.

Definition 4.3.13. *Let $\mathcal{C}$ be a topological prismatoid with bases B^+ and B^-. A* prismatoid shelling of $\mathcal{C}$ from B^+ to B^- *is an ordering $F_1, \ldots, F_K$ of the facets of $\mathcal{C}$ with the following property: for each $i = 1, \ldots, K$, the intersection of $|F_i|$ with $|B^+| \cup |F_1|, \cup \cdots \cup |F_{i-1}|$ is a $(d-2)$-ball in the boundary complex of F_i. Here, the notation $|\cdot|$ applied to a subset of vertices means the subcomplex induced by them.*

Proposition 4.3.14. *Shellability is preserved under the "strong d-step construction" of Theorem 4.3.5.*

Proof. Let $\mathcal{C}$ be a prismatoid with polytopal bases and let $\mathcal{C}'$ be the prismatoid obtained from it in the proof of Theorem 4.3.5. Let $F_1, F_2, \ldots,$ is a shelling order of $\mathcal{C}$.

By the way the sphere S^+ is constructed in that proof, there is a line shelling of S^+ that completely shells the half-sphere S_1^+ first, and then the half-sphere S_2^+. The shelling of $\mathcal{C}'$ is then as follows: Following the shelling order of S^+, shell first $S_1^+ * v_1$, then $S_2^+ * v_2$. After that is done, do $F_1 * v_1, F_1 * v_2, F_2 * v_1, F_2 * v_2, \ldots,$. □

Remark 4.3.15. *It is not clear to us whether the bases of a prismatoid that is shellable in the sense of Definition 4.3.13 have to be shellable themselves.*

4.4 Metaheuristics and implementation

In this section, we show our approach to find non-*d*-step topological prismatoids with few vertices.

The general idea is to start with a 28 vertices prismatoid as defined in [55], and perform flips on it attempting to remove vertices while preserving its width. A general, well known framework to do this is *simulated annealing*. We take this one, instead of the smaller one with 25 vertices also constructed in [55], as a starting point because it has much more symmetry.

Simulated annealing is a very common metaheuristic algorithm for optimization problems, used when we have a search space and an "adjacency relation" between pairs of feasible solutions. The idea is to perform a random walk through the state graph of a problem, but favoring moves that improve the desired objective function over moves that do not. It has been used successfully in combinatorial topology to simplify simplicial complexes while preserving a condition (typically their homeomorphism type) [9] or, in conjunction with other strategies, to tackle the problem of sphere recognition [39, 38].

There is a variable, the *temperature*, regulating the probability assigned to each possible step as a function of how much it improves or worsens the objective. At higher temperatures the choice is more random; when the system cools down it converges to accepting only improving moves. Formally, the probability of accepting a step that increases cost by Δc at temperature t is:

$$\begin{cases} 1 & \text{if } \Delta c < 0, \\ \exp(-\Delta c/t) & \text{if } \Delta c \geq 0. \end{cases}$$

Note that it is very important to choose the potential step with uniform probability, among all the neighbor of the current state. This "a priori" probability distribution, together with the cooling schedule, produces an "a posteriori" probability distribution of performing the step. This gives higher probability to improving steps, but also gives chance to worsening steps at high temperatures.

That is, areas of the graph with smaller values of the objective function are more likely to be explored. As the temperature cools down, the random walk will focus on these areas and make optimizations with more detail. Loosely speaking, the first few iterations of the algorithm are more exploration-focused, and the last iterations are exploitation-focused.

Formally, simulated annealing requires the following aspects to be decided:

- A *state graph* representing the feasible states of the problem and an adjacency relation, plus an initial state.
- A *cooling schedule*, that defines temperature as a function of time, thus modulating the probability of acceptance of a cost-increasing step.
- An appropriate *objective function* that we aim to minimize.

In our particular problem, our *state graph* consists of all non-d-step topological 4-prismatoids, with an edge between a pair of prismatoids if they differ by a flip. This graph is undirected, since every flip is reversible by another flip.

There is a lot of research on *cooling schedules* for different problems. It is known that SA converges to the global optimum for a certain cooling schedule [8], but it is too slow for any practical application. Since the best schedule depends on the problem, several adaptive schedules have been proposed too. However, the most common approach is to define a geometric cooling schedule, of the form $T_t = t_0 * e^{st}$, where the parameters t_0 (initial temperature), s (cooling speed) and the number of iterations are adjusted manually. Since the flipping operation is very fast, we have chosen a slow schedule with a high number of iterations.

The particular parameters have been obtained by trial and error. We used cooling schedule $T(k) = 1000 \cdot 0.99997^k$ and 500000 iterations for each run.

4.4.1 The objective function

The objective function guides our algorithm towards non-d-step topological prismatoids with few vertices. This is, prismatoids with many vertices should have higher cost.

A naive approach would be to just take the number of vertices as an objective function. But this objective function has large *plateaus*, connected subgraphs of the state graph with constant number of vertices, and it does not push our state towards less vertices. So we have to find a way to break ties between prismatoids with the same number of vertices by giving less cost to prismatoids from which we expect it to be easier to remove vertices. That is, the objective function we want has the number of vertices as a main component, plus a smaller (heuristic) *tie-breaker* that pushes the algorithm in the right direction.

A flip that removes a vertex must be of type $(1, d)$; therefore, in order to perform it we must have a vertex with exactly $d + 1$ neighbors, which is the smallest possible size for the neighborhood of a vertex in a prismatoid. So, a good approach is to try to reduce the size of the neighborhood of a vertex until it is precisely $d + 1$.

However, just taking the size of the smallest neighborhood as a tie-breaker is not sensitive enough because the algorithm will then not try to reduce the neighborhoods of other vertices. For this reason, we use as a tie-breaker a generalized mean of the sizes of neighborhoods of vertices. More precisely, the objective function that achieved the best performance in practice among the ones we tried is

$$\operatorname{cost}(\mathcal{C})) = |V(\mathcal{C})| + \varepsilon \left(\frac{\sum_{v \in V(\mathcal{C})} |\operatorname{neigh}(v)|^{-3}}{|V(\mathcal{C})|} \right)^{-1/3} .$$

4.4.2 Data structures

A proper "topological prismatoid" data structure for our problem needs to allow for the following operations:

- Construction from the list of facets.
- Check if a set of vertices is a face.
- Iterate through the maximal subfaces of a face.
- Iterate through the minimal superfaces of a face.
- Perform a flip.
- Compute the width of the prismatoid.
- Get a valid random flip (with uniform probability).

We implement a prismatoid $\mathcal{C}$ as a map of pairs (face,neighborhood), indexed by the faces in $\mathcal{C}$. Faces and neighborhoods themselves are of type "set of integers". We also store the bases of $\mathcal{C}$ in the same manner.

Observe this implicitly gives us the Hasse diagram (maximal subfaces and minimal superfaces): Each face F is directly above those of the form $F \setminus \{v\}$ for $v \in F$, and directly below those of the form $F \cup \{w\}$ for $w \in \operatorname{neigh}(F)$. There is no need to store the adjacency graph, because it is implicit in the Hasse diagram. The facets adjacent to a facet F are computed from the neighborhood of the non-boundary ridges in F. Boundary and interior ridges are distinguished by their neighborhoods having $d+1$ and d elements, respectively.

To compute and update the width we store, for each facet, the distance to the first base (chosen arbitrarily but once and for all) and the number of paths achieving that distance. In this way we do not need to explore again all facets to compute the new width after performing a flip, we just need to update the values that change. For this, when we perform the flip the new facets are inserted into a queue, and the distances are updated by cascading through the prismatoid.

A flip is implemented simply by removing the old faces and inserting the new inserted faces. What is not so straightforward, and needs to be addressed, is how to implement an unbiased generation of random flips among all the possible ones.

For this, we imitate to some extent the technique used in polymake [27]. In polymake, there is a set of pairs (f, l), called "options" satisfying some conditions for flipability, in particular conditions (1) and (2) of Section 4.3.1. The flips are categorized by dimension of f. But the list of candidate pairs (f, l) is very hard to update after a flip is performed. Among other things, some potential flips may change their f and l while preserving their support $f \cup l$.

Since the support of every flip is the neighborhood of a ridge, one could simplify this by using the list of ridges instead of the pairs (f, l) as input to generate a random flip. But choosing randomly from the list of ridges creates bias: some flips are more likely than others since several ridges (actually $|l|$ of them) correspond to the same flip.

To avoid these drawbacks we store and update the list of *ridge-neighborhoods*. That is, for each ridge F we store the facet F_1 containing F or the union $F_1 \cup F_2$ of the two facets containing F depending on whether F_1 is in the boundary or the interior. This is very easy to update, and it also makes it very easy to spot vertex-adding flips (which correspond to boundary ridge-neighborhoods and are characterized by having d instead of $d + 1$ elements). Since there is a bijection between flip-defining ridge-neighborhoods and flips, via the (f, l, v) formalism introduced in Section 4.3.1, it is easy to choose flips uniformly at random: choose a random ridge-neighborhood and discard non-valid ones.

We find this approach more stable and requiring less changes to the data structure than the ones based on (f, l) pairs or in ridges alone.

4.5 Results. Small non-d-step prismatoids and spheres

As said in the previous section, we ran our algorithm with cooling schedule $T(k) = 1000 \cdot 0.99997^k$ and 500000 iterations for each run. We let it run for three days on an openSuse 42.3 Linux machine with 16 GB of RAM and an AMD Phenom X6 1090T processor, after which we had concluded 4093 runs. We thus obtained 4093 non-d-step topological 4-prismatoids, with number of vertices ranging between 14 and 28. Figure 4.3 shows the distribution we obtained for the number of vertices alone (top) and for number of vertices versus number of facets (bottom).

For the rest of this section we focus on the 4 smallest examples, with 14 vertices. We call them #1039, #1963, #2669 and #3513 since these are their indices among the 4093 experiments that we did. They are listed in Tables 4.1–4.4. Vertices from one base are labeled 0 to 6 and vertices from the other are a to g. In the tables, the facets of each prismatoid are grouped by layers, where a layer consists of all facets sharing the number of vertices they have from each base. It is remarkable that the four examples obtained have a lot of similarities:

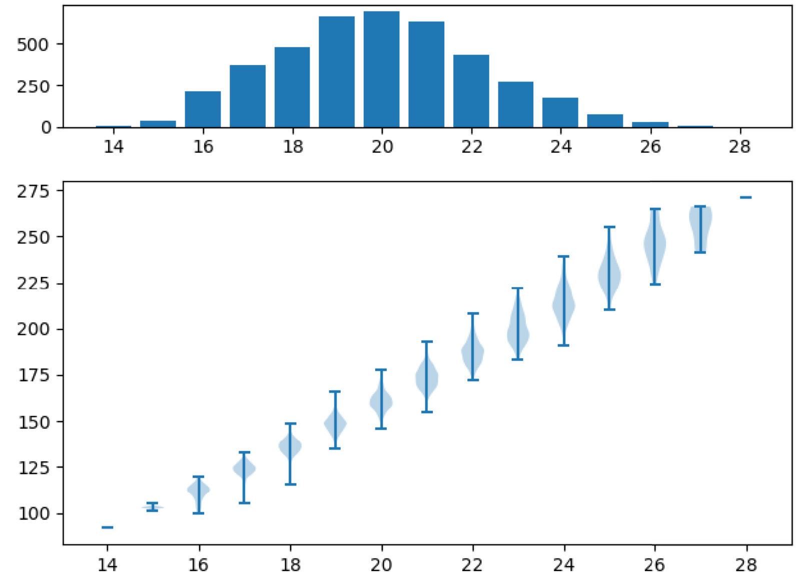

Figure 4.3: Top: distribution of the 4093 prismatoids by number of vertices. Bottom: distribution by number of vertices and facets. The initial prismatoid has 28 vertices and 272 facets.

- They have combinatorially isomorphic bases. Indeed, in all cases the list of facets of the bases are as follows. (To relate this to the tables, observe that the bases correspond to the first and last layer in the prismatoid, removing from each facet the unique vertex from the other base).

$$\begin{array}{ccc} & 0123 & \\ 0134 & 0234 & 1234 \\ 0145 & 0245 & 1245 \\ 0156 & 0256 & 1256 \\ & 0126 & \end{array} \qquad \text{and} \qquad \begin{array}{ccc} & abcd & \\ abde & acde & bcde \\ abef & acef & bcef \\ abfg & acfg & bcfg \\ & abcg. & \end{array}$$

  Observe that both are the face complex of the stacked 4-polytope with seven vertices. That is, they are the boundaries of the stacked 4-balls $\{01234, 01245, 01256\}$ and $\{abcde, abcef, abcfg\}$, respectively.

- They have the same f-vector $(14, 85, 220, 241, 92)$.

- They are all shellable, with a shelling that is monotone on layers. That is, no facet of one layer is used until finishing the previous layer. In the tables, facets within each layer are given in a shelling order.

- The vector of number of facets in different layers is the same $(11, 35, 35, 11)$, and symmetric, in three of them. In #2669 we get the slightly asymmetric vector $(11, 34, 36, 11)$.

- Their reduced incidence patterns, shown in Figure 4.4, are very similar. In #1963 and #3513 the reduced incidence patterns coincide with the one in the right part

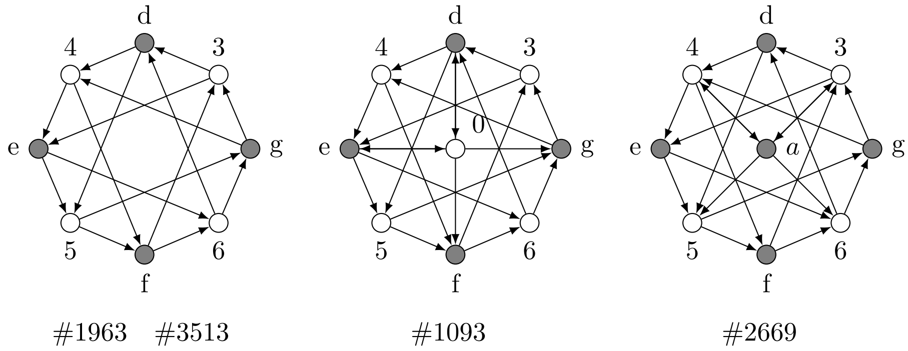

Figure 4.4: The reduced incidence patterns of the four smallest non-d-step 4-prismatoids

of Figure 4.2, minimal by Lemma 4.3.12. In #1039 and #2669 they are are almost the same, except each of them has a single "outlier" facet incident to a base ($0bcde$ in #1039 and $0234a$ in #2669) that introduces a new vertex in the pattern and creates directed cycles of length two. In particular, in these two Proposition 4.3.10 is not enough to prove the non-d-step property.

Note that the starting prismatoid of the algorithm had *the other* reduced incidence pattern of minimal size, the one in the left in Figure 4.2.

We do not know whether any of the four is polytopal.

0256g	126cg	123ae	025af	14abf	26abf	25abf	5acde
0245f	015bg	124af	025bg	13abf	26ace	25abe	5abde
1256g	015bf	123af	024af	13acd	26abe	01bcg	3bcfg
1245f	014ae	123bf	024ae	13acg	05abf	06bcg	3acfg
0234e	013ae	123bg	025ae	13abg	03ace	26bcg	3abfg
0123d	013ad	123cd	025be	06bce	05ace	23bcg	0bcde
1234e	016cd	123cg	125bg	04abd	03acd	23bcf	6bcef
0126d	016cg	026bg	125bf	04bcd	05acd	23acf	6acef
0156g	014bf	026cd	023ce	14bcg	05cde	23ace	6abef
0145f	014ad	026ce	023cd	14abg	05bde	26bcf	4abcg
0134e	014cd	026be	124ae	14acg	05abd	26acf	4abcd
	014bc	126cd		14acd	04abf		

Table 4.1: Prismatoid #1039

0126d	014ae	023bd	126cg	14abe	24acd	25bcd	5abde
0123d	024ae	123bd	125cg	24abe	12acf	25bce	6abef
0134e	013ae	023be	125cf	24abd	25acf	26bce	3abfg
0234e	014af	026be	015cf	25abe	25acd	13acg	5acde
1234e	024af	123be	013cf	25abd	05acf	14acg	6acef
0145f	025af	124be	015cg	13abe	05ace	16bcg	4abcd
0245f	025ae	124bd	025cg	03abe	06ace	06bcg	5bcde
1245f	124af	016bg	025ce	06abe	06acf	06bcf	6bcef
0156g	013af	013bg	026cg	03abf	13acf	03bcf	3acfg
0256g	016bd	126cd	026ce	06abf	16bcd	14bcg	4abcg
1256g	026bd	124cd	013cg	13abg	14bcd	03bcg	3bcfg
	013bd	124ac		14abg	26bcd		

Table 4.2: Prismatoid #1963

0156g	023ad	126cg	024af	24abg	06bcg	04abf	5abde
0256g	013ad	015cg	026af	23abg	26bcg	04abe	5acde
1256g	026ad	016cd	026bf	13acg	23bcg	06abe	5bcde
0123d	013ae	126cd	025bf	14acg	05bcg	14abe	4abcg
0126d	034ae	123cd	125bf	13acd	15bcg	14abd	4abcd
0134e	234ae	123cg	124bf	23acd	14bcg	15abd	3abfg
1234e	123ae	123ag	015bf	26acd	14bcd	15abe	6abef
0234a	124ae	124ag	015be	06acd	15bcd	06abf	3acfg
0145f	026bg	124bg	014bf	14acd	05bce	23acf	6acef
0245f	025bg	015cd	014be	05acd	06bce	26acf	3bcfg
1245f	125bg	015ad		05ace	24abf	23bcf	6bcef
	016cg	015ae		06ace	23abf	26bcf	

Table 4.3: Prismatoid #2669

0156g	015ag	014bg	014ce	04abg	14bcg	04bcd	5abde
0256g	025bg	024af	013ce	05abg	14acg	16bcd	3abfg
1256g	125bg	024ae	023ce	15abg	14ace	26bcd	6abef
0134e	026bg	124af	025ce	05abd	04ace	25bcd	4abcg
0234e	126bg	124ae	025cd	04abd	05ace	25bce	4abcd
1234e	016bg	123ae	023cd	15abf	13ace	26bce	5acde
0126d	016bd	025af	013cd	25abf	13acg	26ace	5bcde
0123d	026bd	025ae	123cd	25abe	13bcg	26acf	3acfg
0145f	025bd	123af	126cd	26abf	23ace	23acf	3bcfg
0245f	015af	125bf	126cf	26abe	05acd	16bcf	6acef
1245f	014af	126bf	123cf	13abf	04acd	13bcf	6bcef
	014ag	014bc		13abg	01bcd		

Table 4.4: Prismatoid #3513

Bibliography

[1] Karim A Adiprasito and Bruno Benedetti, *The Hirsch conjecture holds for normal flag complexes*, Mathematics of Operations Research **39** (2014), no. 4, 1340–1348.

[2] Marianne Akian, Stéphane Gaubert, and Alexander Guterman, *Tropical polyhedra are equivalent to mean payoff games*, International Journal of Algebra and Computation **22** (2012), no. 01, 1250001.

[3] Daniele Alessandrini, *Logarithmic limit sets of real semi-algebraic sets*, Adv. Geom. **13** (2013), no. 1, 155–190.

[4] Xavier Allamigeon, Pascal Benchimol, Stéphane Gaubert, and Michael Joswig, *Combinatorial simplex algorithms can solve mean payoff games*, SIAM Journal on Optimization **24** (2014), no. 4, 2096–2117.

[5] ______, *Log-barrier interior point methods are not strongly polynomial*, SIAM Journal on Applied Algebra and Geometry **2** (2018), no. 1, 140–178.

[6] Noga Alon and Joel H. Spencer, *The probabilistic method*, fourth ed., Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2016.

[7] Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee, *Voronoi diagrams and Delaunay triangulations*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2013.

[8] Dimitris Bertsimas, John Tsitsiklis, et al., *Simulated annealing*, Statistical science **8** (1993), no. 1, 10–15.

[9] Anders Björner and Frank H Lutz, *Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the poincaré homology 3-sphere*, Experimental Mathematics **9** (2000), no. 2, 275–289.

[10] Roswitha Blind and Peter Mani-Levitska, *Puzzles and polytope isomorphisms*, Aequationes mathematicae **34** (1987), no. 2-3, 287–297.

[11] Jesús Burgueño, *Geografía política de la españa constitucional: la división provincial*, Centro de estudios constitucionales Madrid, 1996.

[12] L. Caccetta and W. F. Smyth, *Graphs of maximum diameter*, Discrete Math. **102** (1992), no. 2, 121–141.

[13] L. Paul Chew and Robert L. Scot Dyrsdale III, *Voronoi diagrams based on convex distance functions*, Proceedings of the first annual symposium on Computational geometry, ACM, 1985, pp. 235–244.

[14] Francisco Criado, Michael Joswig, and Francisco Santos, *Tropical bisectors and voronoi diagrams*, arXiv preprint arXiv:1906.10950 (2019).

[15] Francisco Criado and Andrew Newman, *Randomized construction of complexes with large diameter*, Discrete & Computational Geometry (2020).

[16] Francisco Criado and Francisco Santos, *The maximum diameter of pure simplicial complexes and pseudo-manifolds*, Discrete & Computational Geometry **58** (2017), no. 3, 643–649.

[17] ———, *Topological prismatoids and small simplicial spheres of large diameter*, Experimental Mathematics (2019), 1–13.

[18] George B Dantzig and Mukund N Thapa, *Linear programming 1: introduction*, Springer Science & Business Media, 2006.

[19] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars, *Computational geometry*, third ed., Springer-Verlag, Berlin, 2008, Algorithms and applications.

[20] Jules Depersin, Stéphane Gaubert, and Michael Joswig, *A tropical isoperimetric inequality*, Sém. Lothar. Combin. – FPSAC 2017 – **78B** (2017), Art. 27, 12.

[21] Mike Develin, *The moduli space of n tropically collinear points in $\mathbb{R}^d$*, Collect. Math. **56** (2005), no. 1, 1–19.

[22] Jack Edmonds, *Systems of distinct representatives and linear algebra*, J. Res. Nat. Bur. Standards Sect. B **71** (1967), no. 4, 241–245.

[23] P. Erdős and L. Lovász, *Problems and results on 3-chromatic hypergraphs and some related questions*, (1975), 609–627. Colloq. Math. Soc. János Bolyai, Vol. 10.

[24] Xin Gui Fang and George Havas, *On the worst-case complexity of integer gaussian elimination*, Proceedings of the 1997 international symposium on Symbolic and algebraic computation, 1997, pp. 28–31.

[25] Steven Fortune, *A sweepline algorithm for Voronoĭ diagrams*, Algorithmica **2** (1987), no. 2, 153–174.

[26] Joseph Fourier, *Histoire de l'académie, partie mathématique (1824)*, Mémoires de l'Académie des sciences de l'Institut de France **7** (1824).

[27] Ewgenij Gawrilow and Michael Joswig, `polymake`*: a framework for analyzing convex polytopes*, Polytopes—combinatorics and computation (Oberwolfach, 1997), DMV Sem., vol. 29, Birkhäuser, Basel, 2000, pp. 43–73.

[28] Chan He, Horst Martini, and Senlin Wu, *On bisectors for convex distance functions*, Extracta Math. **28** (2013), no. 1, 57–76.

[29] Christian Icking, Rolf Klein, Ngọc-Minh Lê, and Lihong Ma, *Convex distance functions in 3-space are different*, Fund. Inform. **22** (1995), no. 4, 331–352, Computational geometry (San Diego, CA, 1993).

[30] Christian Icking, Rolf Klein, Ngoc-Minh Lê, Lihong Ma, and Francisco Santos, *On bisectors for convex distance functions in 3-space*, Proceedings of the 11th Canadian Conference on Computational Geometry, 1999.

[31] Ivan Izmestiev and Michael Joswig, *Branched coverings, triangulations, and 3-manifolds*, Advances in Geometry **3** (27 Mar. 2003), no. 2, 191 – 225.

[32] Ivan Izmestiev, Steven Klee, and Isabella Novik, *Simplicial moves on balanced complexes*, Advances in Mathematics **320** (2017), 82–114.

[33] Philipp Jell, Claus Scheiderer, and Josephine Yu, *Real tropicalization and analytification of semialgebraic sets*, International Mathematics Research Notices (2020).

[34] Michael Joswig, *Tropical halfspaces*, Combinatorial and computational geometry, Math. Sci. Res. Inst. Publ., vol. 52, Cambridge Univ. Press, Cambridge, 2005, pp. 409–431.

[35] ———, *Essentials of tropical combinatorics*, Book in preparation **1** (2014), no. 226, 9.

[36] Michael Joswig and Katja Kulas, *Tropical and ordinary convexity combined*, Adv. Geometry **10** (2010), 333–352.

[37] Michael Joswig and Georg Loho, *Weighted digraphs and tropical cones*, Linear Algebra Appl. **501** (2016), 304–343.

[38] Michael Joswig, Frank H Lutz, and Mimi Tsuruga, *Heuristics for sphere recognition*, International Congress on Mathematical Software, Springer, 2014, pp. 152–159.

[39] ———, *Sphere recognition: Heuristics and examples*, arXiv preprint arXiv:1405.3848 (2015), 152–159.

[40] Michael Joswig and Benjamin Schröter, *The tropical geometry of shortest paths*, arXiv preprint arXiv:1904.01082 (2019).

[41] Gil Kalai, *Many triangulated spheres*, Discrete & Computational Geometry **3** (1988), no. 1, 1–14.

[42] Gil Kalai and Daniel J Kleitman, *A quasi-polynomial bound for the diameter of graphs of polyhedra*, arXiv preprint math/9204233 (1992).

[43] Narendra Karmarkar, *A new polynomial-time algorithm for linear programming*, Proceedings of the sixteenth annual ACM symposium on Theory of computing, 1984, pp. 302–311.

[44] Leonid Genrikhovich Khachiyan, *A polynomial algorithm in linear programming*, Doklady Akademii Nauk, vol. 244, Russian Academy of Sciences, 1979, pp. 1093–1096.

[45] Edward D Kim and Francisco Santos, *An update on the hirsch conjecture*, Jahresbericht der Deutschen Mathematiker-Vereinigung **112** (2010), no. 2, 73–98.

[46] Victor Klee, David W Walkup, et al., *The d-step conjecture for polyhedra of dimension d< 6*, Acta Mathematica **117** (1967), 53–78.

[47] J. J. Kosowsky and Alan L. Yuille, *The invisible hand algorithm: Solving the assignment problem with statistical physics*, Neural Networks **7** (1994), no. 3, 477–490.

[48] Thomas Lam and Alexander Postnikov, *Alcoved polytopes. I*, Discrete & Computational Geometry **38** (2007), no. 3, 453–478.

[49] David G Larman, *Paths on polytopes*, Proceedings of the London Mathematical Society **3** (1970), no. 1, 161–178.

[50] Bo Lin, Anthea Monod, and Ruriko Yoshida, *Tropical foundations for probability & statistics on phylogenetic tree space*, 2018, Preprint `arXiv:1805.12400`.

[51] Georg Loho and Matthias Schymura, *Tropical Ehrhart theory and tropical volume*, arXiv preprint arXiv:1908.07893 (2019).

[52] Diane Maclagan and Bernd Sturmfels, *Introduction to tropical geometry*, Graduate Studies in Mathematics, vol. 161, American Mathematical Society, Providence, RI, 2015.

[53] Peter Mani and David W Walkup, *A 3-sphere counterexample to the wv-path conjecture*, Mathematics of Operations Research **5** (1980), no. 4, 595–598.

[54] Horst Martini and Konrad J Swanepoel, *The geometry of Minkowski spaces, a survey. part ii*, Expositiones mathematicae **22** (2004), no. 2, 93–144.

[55] Benjamin Matschke, Francisco Santos, and Christophe Weibel, *The width of five-dimensional prismatoids*, Proceedings of the London Mathematical Society **110** (2015), no. 3, 647–672.

[56] Grigory Mikhalkin, *Enumerative tropical algebraic geometry in $\mathbb{R}^2$*, J. Amer. Math. Soc. **18** (2005), no. 2, 313–377.

[57] Rolf H Möhring, Martin Skutella, and Frederik Stork, *Scheduling with and/or precedence constraints*, SIAM Journal on Computing **33** (2004), no. 2, 393–415.

[58] J. W. Moon, *On the diameter of a graph*, Michigan Math. J. **12** (1965), 349–351.

[59] Andrew Newman, *Small simplicial complexes with prescribed torsion in homology*, Discrete & Computational Geometry (2017).

[60] ———, *A lower bound on the number of homotopy types of simplicial complexes on n vertices*, arXiv: 1804.06787, 2018.

[61] J. R. Norris, *Markov chains*, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 2, Cambridge University Press, Cambridge, 1998, Reprint of 1997 original.

[62] Udo Pachner, *Pl homeomorphic manifolds are equivalent by elementary shellings*, European journal of Combinatorics **12** (1991), no. 2, 129–145.

[63] Francisco Santos, *A counterexample to the Hirsch conjecture*, Ann. of Math. (2) **176** (2012), no. 1, 383–412.

[64] ———, *Recent progress on the combinatorial diameter of polytopes and simplicial complexes*, Top **21** (2013), no. 3, 426–460.

[65] Steve Smale, *Mathematical problems for the next century*, The mathematical intelligencer **20** (1998), no. 2, 7–15.

[66] David Speyer and Bernd Sturmfels, *The tropical Grassmannian*, Adv. Geom. **4** (2004), no. 3, 389–411.

[67] Rade T. Živaljević, *Topological methods in discrete geometry*, Handbook of Discrete and Computational Geometry (Csaba D. Tóth, Jacob E. Goodmann, and Joseph O'Rourke, eds.), CRC Press, 2018, 3rd edition.

[68] Peter Widmayer, Ying-Fung Wu, and Chak-Kuen Wong, *On some distance problems in fixed orientations*, SIAM Journal on Computing **16** (1987), no. 4, 728–746.

[69] Boris Yamnitsky and Leonid A Levin, *An old linear programming algorithm runs in polynomial time*, 23rd Annual Symposium on Foundations of Computer Science (sfcs 1982), IEEE, 1982, pp. 327–328.

[70] Günter M Ziegler, *Lectures on polytopes*, vol. 152, Springer Science & Business Media, 2012.