

Chapter 10

Convex Polytopes

Recall from Definition 5.8 that a **convex polytope** is the convex hull of a finite point set $P \subset \mathbb{R}^d$. In this chapter, we take a closer look at their structures and reveal their links to high-dimensional Delaunay triangulations and Voronoi diagrams. For convenience, we shall omit the attribute *convex* and refer to them simply as *polytopes*. In the sequel we will be borrowing a lot of material from Ziegler's classical book *Lectures on Polytopes* [12], sometimes without proofs as they would take us too far from geometry.

We are already familiar with polytopes in dimension $d = 2$, which are just convex polygons; see Figure 10.1. These are boring in the combinatorial sense: the vertex-edge graph is always a cycle.




Figure 10.1: In \mathbb{R}^2 , convex polytopes are convex polygons.

Polytopes in dimension $d = 3$ are more interesting, as they own a richer combinatorial structure. The most popular examples are the five platonic solids; see Figure 10.2.




Figure 10.2: The five platonic solids. (Images from Wikipedia [3, 2, 10, 1, 4])

Despite this diversity, the vertex-edge graphs of 3-dimensional polytopes are well-understood, due to the following classical result. (See Lecture 4 in Ziegler’s book [12] for a thorough treatment.)

Theorem 10.1 (Steinitz). *A graph G is the vertex-edge graph of a 3-dimensional polytope if and only if G is planar and 3-connected.*

We have already encountered 3-connected planar graphs in Chapter 2. Recall that Whitney’s Theorem 2.27 states that every such graph has a unique combinatorial embedding in the plane. Here, Steinitz’s theorem says that it also admits a geometric embedding as the skeleton of some polytope in \mathbb{R}^3 . One can easily verify the theorem on the five platonic solids; for example, Figure 10.3 shows the vertex-edge graph of the octahedron, which is clearly planar and 3-connected.




Figure 10.3: *The vertex-edge graph of the octahedron*

The theorem implies that a polytope in \mathbb{R}^3 with n vertices has at most $3n - 6$ edges and $2n - 4$ faces, by Corollary 2.5. What happens in higher dimensions? In particular, we want to understand how complicated a polytope in \mathbb{R}^d can be. For example, how many edges can a 4-dimensional polytope with n vertices have? Is it still $O(n)$ as for $d = 2, 3$? To discuss this, we first have to define “vertices” and “edges” formally—our intuition unfortunately stops in \mathbb{R}^3 . In fact, it is useful to define the more general notion of *faces* which subsumes vertices and edges.

10.1 Faces of a Polytope

In studying general dimension d , linear algebra tools are prominent. For a quick refresher we refer the reader to Chapter 5. Also recall that the dimension of a linear space is the maximum size of its linear independent subset. The dimension of an affine space is the maximum size of its affinely independent subset.

Let $\mathcal{P} = \text{conv}(P)$ be a polytope. Its **dimension** $\dim(\mathcal{P})$ is the dimension of its affine hull. The polytope is **full-dimensional** if $\dim(\mathcal{P}) = d$. Many results are stated for full-dimensional polytopes only, but this is not really a restriction: If $\dim(\mathcal{P}) < d$ then we can study it in the affine subspace $\text{aff}(\mathcal{P}) \cong \mathbb{R}^{\dim(\mathcal{P})}$ where \mathcal{P} becomes full-dimensional.

Definition 10.2. Let $\mathcal{P} \subset \mathbb{R}^d$ be a polytope. We call $F \subseteq \mathcal{P}$ a **face** of \mathcal{P} if there is a hyperplane

$$h = \left\{ x \in \mathbb{R}^d : \sum_{i=1}^d h_i x_i = h_{d+1} \right\}$$

such that $F = \mathcal{P} \cap h$ and $\mathcal{P} \subset h^+$, where

$$h^+ = \left\{ x \in \mathbb{R}^d : \sum_{i=1}^d h_i x_i \geq h_{d+1} \right\}$$

is the closed positive halfspace bounded by h .¹ We say that the h **supports** face F .

You should think of a face as the intersection of \mathcal{P} with a hyperplane “tangent” to \mathcal{P} . Figure 10.4 illustrates this notion. The **dimension** of a face is the dimension of its affine hull. A face of dimension k is called a k -face. Conventionally,

- 0-faces are called *vertices*,
- 1-faces are called *edges*,
- $(\dim(\mathcal{P}) - 2)$ -faces are called *ridges*, and
- $(\dim(\mathcal{P}) - 1)$ -faces are called *facets*.

For example, the octahedron in Figure 10.2(c) has 8 facets, 12 edges (which are also ridges), and 6 vertices. The dodecahedron in Figure 10.2(d) has 12 facets, 30 edges(=ridges), and 20 vertices.




Figure 10.4: Two faces (an edge and a vertex) with supporting hyperplanes.

Degeneracy occurs if we set $h_1 = \dots = h_d = 0$ in the definition.² If $h_{d+1} = 0$ then $h = h^+ = \mathbb{R}^d$, so this hyperplane supports \mathcal{P} . If $h_{d+1} < 0$ then $h = \emptyset, h^+ = \mathbb{R}^d$, so this hyperplane supports \emptyset . These two are called degenerate faces of \mathcal{P} ; the others are called proper faces.

¹ Note that an inequality $\sum_{i=1}^d h_i x_i \leq h_{d+1}$ is equivalent to $\sum_{i=1}^d (-h_i) x_i \geq -h_{d+1}$, so sticking to positive halfspaces in the definition is no loss of generality.

² In Section 1.2 we did not allow such, but here we need it.

Exercise 10.3. Show that every facet of a full-dimensional polytope has a unique supporting hyperplane: its affine hull.

The definition of a face, in its current form, is cumbersome to work with. In particular, to verify a supporting hyperplane h , we have to reason about its interaction with the continuous mass \mathcal{P} . Thankfully, the lemma below reduces the verification to a finite set.

Lemma 10.4. Let $\mathcal{P} = \text{conv}(P) \subset \mathbb{R}^d$ be a polytope. If a hyperplane h satisfies $P \subset h^+$, then $\mathcal{P} \subset h^+$ and $\mathcal{P} \cap h = \text{conv}(P \cap h)$.

To get the intuition for the lemma as well as its proof, let us rephrase it: Imagine a hyperplane h and several points, some sitting on h while others living strictly on one side of h . Then the convex hull of the points should also dwell in that side and, moreover, intersect h in the zone that “fills between” the points on h .

Proof. As $P \subset h^+$ and h^+ is convex, the first claim follows immediately. The intersection $\mathcal{P} \cap h$ is convex as both \mathcal{P} and h are convex, therefore $\text{conv}(P \cap h) \subseteq \mathcal{P} \cap h$.

It remains to show $\mathcal{P} \cap h \subseteq \text{conv}(P \cap h)$. Assume $h = \{x \in \mathbb{R}^d : \sum_{i=1}^d h_i x_i = h_{d+1}\}$. Let $p \in (\mathcal{P} \setminus P) \cap h$. Since $p \in \mathcal{P}$ we can express $p = \sum_{q \in Q} \lambda_q q$ as a convex combination of some other points $Q \subseteq P$, where $\lambda_q > 0$ for all $q \in Q$. Since $p \in h$ we know

$$h_{d+1} = \sum_{i=1}^d h_i p_i = \sum_{i=1}^d h_i \sum_{q \in Q} \lambda_q q_i = \sum_{q \in Q} \lambda_q \sum_{i=1}^d h_i q_i \geq \sum_{q \in Q} \lambda_q h_{d+1} = h_{d+1}.$$

where the “ \geq ” uses the fact that $P \subset h^+$. As the two ends are equal, the inequality is in fact an equality. But recall $\lambda_q > 0$ for all $q \in Q$, so we must have $\sum_{i=1}^d h_i q_i = h_{d+1}$ for all $q \in Q$. In other words, $Q \subseteq P \cap h$. Therefore, $p \in \text{conv}(Q) \subseteq \text{conv}(P \cap h)$. \square

As an immediate consequence, every face of $\text{conv}(P)$ is a convex hull of some points in P . In particular, there can be at most $2^{|P|}$ faces—a finite number as one would expect, but not at all obvious from the definition!

Lemma 10.4 is central to many proofs in this chapter; let us see an application right away. You might have speculated that the extreme points of a set P (cf. Definition 5.8) coincide with the vertices of polytope $\text{conv}(P)$. Indeed this is true, up to the formal subtlety that a vertex is a singleton set rather than a point (we will later ignore this nuance, but it is good to have talked about it once).

Lemma 10.5. Let $\mathcal{P} = \text{conv}(P) \subset \mathbb{R}^d$ be a polytope. Then p is an extreme point of \mathcal{P} if and only if $\{p\}$ is vertex of \mathcal{P} .

Proof. If $p = (p_1, \dots, p_d)$ is an extreme point of P , then the compact convex sets $\{p\}$ and $\text{conv}(P \setminus \{p\})$ are disjoint. By the Separation Theorem 5.19, there is a hyperplane h that strictly separates them. In formulas, there exist non-degenerate hyperplane parameters $h_1, \dots, h_{d+1} \in \mathbb{R}$ such that

$$\sum_{i=1}^d h_i p_i < h_{d+1} \quad \text{and} \quad \sum_{i=1}^d h_i q_i > h_{d+1} \quad \forall q \in \text{conv}(P \setminus \{p\}).$$

Let us decrease h_{d+1} until the first inequality becomes tight. At that moment we have in particular

$$\sum_{i=1}^d h_i p_i = h_{d+1} \quad \text{and} \quad \sum_{i=1}^d h_i q_i > h_{d+1} \quad \forall q \in P \setminus \{p\},$$

meaning that $P \subset h^+$ and $P \cap h = \{p\}$. Applying Lemma 10.4 we obtain $\mathcal{P} \subseteq h^+$ and $\mathcal{P} \cap h = \{p\}$, so the hyperplane h supports $\{p\}$.

For the other direction, let p be a vertex supported by some hyperplane h . Consider the set $P' := P \setminus \{p\}$. Clearly $P' \subset h^+$ and $P' \cap h = \emptyset$. Applying Lemma 10.4 on P' , we derive $\text{conv}(P') \cap h = \emptyset$. But $p \in h$, so $p \notin \text{conv}(P')$, namely p is an extreme point. \square

By $V(\mathcal{P})$ we denote the set of vertices of a polytope \mathcal{P} . Then Proposition 5.10 with Lemma 10.5 imply the following:

Corollary 10.6. $\mathcal{P} = \text{conv}(V(\mathcal{P}))$; moreover, $V(\mathcal{P}) = \bigcap_{P: \text{conv}(P)=\mathcal{P}} P$.

In words, $V(\mathcal{P})$ is the minimal description of a polytope \mathcal{P} as a convex hull of points. With little extra work we can relate vertices with arbitrary faces.

Lemma 10.7. Every face F of a polytope \mathcal{P} is a polytope itself with $V(F) = V(\mathcal{P}) \cap F$.

Proof. Let F be a face supported by some hyperplane h . By Lemma 10.4, $F = \mathcal{P} \cap h = \text{conv}(V(\mathcal{P}) \cap h) = \text{conv}(V(\mathcal{P}) \cap F)$. So F is a polytope. Moreover, $V(F) \subseteq V(\mathcal{P}) \cap F$ due to Corollary 10.6. The converse inclusion is clear, as any hyperplane supporting a vertex $v \in V(\mathcal{P}) \cap F$ in the polytope \mathcal{P} is also supporting v in the face F . \square

Let us consider some concrete examples. Each facet of the octahedron is a triangle (which is a polytope); its three vertices are exactly the vertices of the octahedron that lie on the triangle. Similarly, each edge is a line segment (which is a polytope as well); its two vertices are those of the octahedron that lie on the segment.

In view of Corollary 10.6 and Lemma 10.7, every face F can be encoded by $V(\mathcal{P}) \cap F$ and later restored by taking convex hull. Namely, $F \mapsto V(\mathcal{P}) \cap F$ is an injection from faces of \mathcal{P} to subsets of $V(\mathcal{P})$. In particular, if $|V(\mathcal{P})| = n$ then \mathcal{P} has at most 2^n faces.

Exercise 10.8. Let \mathcal{P} be a polytope with n vertices. Show that \mathcal{P} has at most $\binom{n}{k+1}$ many k -faces, for every $0 \leq k < \dim(\mathcal{P})$.

Specializing for $k = 2$, a polytope with n vertices can have at most $\binom{n}{2}$ edges which doesn't surprise us: the vertex-edge graph cannot be more than complete. For $d = 2, 3$ this is a gross overestimate as we know there can be only $O(n)$ many edges; nevertheless, we can use Exercise 10.8 to upper bound the total number of proper faces by

$$\sum_{k=0}^{\dim(\mathcal{P})-1} \binom{n}{k+1} = O(n^{\dim(\mathcal{P})}),$$

which substantially improves the previous bound 2^n (for $n \rightarrow \infty$ and constant d).

Lemma 10.9. *Let F, G be two faces of a polytope \mathcal{P} . Then $F \cap G$ is also a face of \mathcal{P} . It has vertex set $V(F \cap G) = V(F) \cap V(G)$.*

Proof Sketch. Assume that F and G are supported by hyperplanes $\sum_{i=1}^d a_i x_i = a_{d+1}$ and $\sum_{i=1}^d b_i x_i = b_{d+1}$, respectively. Consider their “mixture”

$$h := \left\{ x \in \mathbb{R}^d : \sum_{i=1}^d (a_i + b_i)x_i = a_{d+1} + b_{d+1} \right\}.$$

It is not hard to check that

- every point $p \in F \cap G$ is lying on h ;
- every point $p \in V(\mathcal{P}) \setminus (V(F) \cap V(G))$ is strictly contained in h^+ .

It follows from Lemma 10.4 that $\mathcal{P} \subset h^+$ and $\mathcal{P} \cap h = \text{conv}(F \cap G) = F \cap G$. So $F \cap G$ is a face supported by h . By Lemma 10.7, its vertex set is exactly $\mathcal{P} \cap (F \cap G) = V(F) \cap V(G)$. \square

Exercise 10.10. *Show that every ridge is incident to exactly two facets.*

For polytopes in \mathbb{R}^3 , Euler’s formula gives us a relation between the number of vertices, edges and facets. In higher dimension this is generalized by the Euler-Poincaré formula. Let us denote by f_k the number of k -faces of a polytope \mathcal{P} .

Theorem 10.11 (Euler-Poincaré formula). *For every d -dimensional polytope we have*

$$\sum_{k=0}^{d-1} (-1)^k f_k = 1 - (-1)^d.$$

When specializing to $d = 3$ we recover the familiar $f_0 - f_1 + f_2 = 2$. We will see an elegant proof of the formula later in Chapter 11; but now let us explore one of its consequences:

Exercise 10.12. *Let $P \subset \mathbb{R}^4$ be a finite set of points in general position and let \mathcal{P} be the polytope defined by the convex hull of P . Show that $f_3 \geq f_0$.*

10.2 The Main Theorem

We already know from Theorem 5.22 that a polytope can be written as the intersection of infinitely many halfspaces. But it seems that most of them are redundant; at least in dimension $d \leq 3$ finitely many halfspaces suffice. Is it true for higher dimensions? This motivates the following definition.

Definition 10.13. *A polyhedron is the intersection of finitely many halfspaces in \mathbb{R}^d .*




Figure 10.5: An (unbounded) polyhedron in \mathbb{R}^2 (intersection of four halfspaces)

Unlike polytopes, polyhedra may be unbounded. For example, the whole space \mathbb{R}^d is a polyhedron (the intersection of no halfspaces), and every halfspace is also a polyhedron (the intersection of one halfspace). Figure 10.5 gives another example in \mathbb{R}^2 .

The faces of a polyhedron \mathcal{P} can be defined in the same way as for polytopes: F is a face if there exists a hyperplane h such that $F = \mathcal{P} \cap h$ and $\mathcal{P} \subset h^+$. For example, the polyhedron in Figure 10.5 has 3 vertices and 4 edges (= facets), two of which unbounded.

By extrapolating from the case $d = 2$ (which is always a bit dangerous, but let's try anyway), it seems that the only thing that can stop a polyhedron from being a polytope is its unboundedness. Conversely, it seems that a polytope is always a (bounded) polyhedron. These are indeed true in any dimension! So polytopes and bounded polyhedra is the same object. This is arguably the most fundamental result in polytope theory, and for this reason, Ziegler calls it the *Main Theorem* [12, Theorem 1.1].

Theorem 10.14 (Main Theorem). *A subset $\mathcal{P} \subset \mathbb{R}^d$ is the convex hull of a finite set of points if and only if \mathcal{P} is a bounded intersection of finitely many halfspaces.*

People usually use the attributes \mathcal{V} -polytope and \mathcal{H} -polytope to mean a polytope represented as a convex hull of points or an intersection of hyperspaces, respectively.

Exercise 10.15. *Let $\mathcal{P} = \bigcap_{i=1}^m h_i^+$ be a full-dimensional polytope, represented as the intersection of m halfspaces h_1^+, \dots, h_m^+ . Prove that each facet of \mathcal{P} is supported by one of the m hyperplanes h_i . (As a hyperplane can by definition support only one facet, \mathcal{P} has at most m facets.)*

It can also be shown [12, Theorem 2.15 (7)] that hyperplanes not supporting a facet are redundant, meaning that we can always write a full-dimensional polytope with m facets in the form $\mathcal{P} = \bigcap_{i=1}^m h_i^+$, where each h_i supports one of the facets. Hence, in the same way that non-extreme points are redundant in representing a \mathcal{V} -polytope, hyperplanes not supporting facets are redundant in representing an \mathcal{H} -polytope.

Corollary 10.16. *Let \mathcal{P} be a full-dimensional polytope. Then every point $p \in \partial\mathcal{P}$ is contained in some facet.*

Proof. Represent $\mathcal{P} = \bigcap_{i=1}^m h_i^+$ as an intersection of facet-supporting hyperplanes. If $p \in \mathcal{P}$ is not contained in any facet, then it is not contained in any of the hyperplanes. So a sufficiently small ball around p would still be in \mathcal{P} , meaning that $p \notin \partial\mathcal{P}$. \square

10.3 Two Examples

Let's look at two families of higher-dimensional polytopes, called *hypercubes* and *simplices*, that naturally generalize the cube and the tetrahedron, respectively. The standard d -dimensional **hypercube** is the set

$$\{x \in \mathbb{R}^d : -1 \leq x_i \leq 1 \quad \forall i \in \{1, \dots, d\}\}.$$

Formally this is a polyhedron, described as the intersection of $2d$ halfspaces. But as the boundedness is clear, the Main Theorem guarantees that it is a polytope. It has at most $2d$ facets by Exercise 10.15; but one can easily show that the number is precisely $2d$ (try to make the argument!). The next exercise is about its vertices.

Exercise 10.17. *Prove that the standard d -dimensional hypercube has 2^d vertices. What are they?*

A d -dimensional simplex, or simply **d -simplex**, is the convex hull of $d + 1$ affinely independent points.

Exercise 10.18. *Prove that any d -simplex has 2^{d+1} faces. Specifically, for every subset Q of its defining points, show that there is a face F with $V(F) = Q$. (This count is maximum possible for polytopes with $d + 1$ vertices, by Lemma 10.7.)*

10.4 Polytope Structure

In this section, we will summarize some more advanced properties of polytopes. All of these classical material can be found in full detail in Ziegler's book [12].

10.4.1 The Graph of a Polytope

For any d -dimensional polytope \mathcal{P} , its vertices and edges form a graph $G(\mathcal{P})$, sometimes also called the 1-skeleton of \mathcal{P} . As we discussed in the beginning of this chapter, these graphs are just cycles for dimension $d = 2$, and triconnected planar graphs for dimension $d = 3$ (Steinitz's Theorem 10.1). It turns out that for higher dimensions d , the graphs are also d -connected, as we will soon show.

Why do we care about these graphs? From a computational viewpoint they are very relevant to *linear programming*, a cornerstone in optimization theory. We will briefly explain the connection here without going into details. In a **linear program** we want to maximize a linear function $c^\top x$ where the variable x is subject to a system of linear inequalities $Ax \leq b$. Each row in $Ax \leq b$ specifies a halfspace, so all the rows together define a polyhedron \mathcal{P} . Let us assume for simplicity that it is non-empty and bounded, hence a polytope by the Main Theorem. Let $\zeta_{\max} := \max_{x \in \mathcal{P}} c^\top x$ be the optimal value of the linear program. Then $c^\top x = \zeta_{\max}$ is a hyperplane whose intersection with \mathcal{P} is the set of optimal solutions. In particular, the set of optimal solutions is a face of \mathcal{P} . Let us

orient every edge $\{v, w\}$ of $G(\mathcal{P})$ so that $v \rightarrow w$ if and only if $c^\top v < c^\top w$. Clearly the oriented graph is acyclic. Further, Proposition 10.19 below implies that every sink is an optimal solution. Thus one way to find an optimal solution is to start from any vertex and follow the directed edges, until we reach a sink. This is the main idea of an entire family of algorithms for linear programming, called the *simplex method*.

Proposition 10.19 (see [12]). *Let \mathcal{P} be a polytope. Orient the graph $G(\mathcal{P})$ as above, according to the linear function $c^\top x$. Show that if vertex $v \in V(\mathcal{P})$ is suboptimal, that is if $c^\top v < \max_{x \in \mathcal{P}} c^\top x$, then there is an edge going out of v .*

In order for the simplex method to work efficiently, the graph $G(\mathcal{P})$ needs to have small diameter. Warren M. Hirsch made the following conjecture in 1957, known as the *Hirsch conjecture*: For any d -dimensional polytope \mathcal{P} with m facets, the diameter of graph $G(\mathcal{P})$ is at most $m - d$.

This conjecture was disproven in 2010 by Francisco Santos [8], who constructed a 43-dimensional polytope with 86 facets whose graph has diameter larger than 43. However, the weaker *polynomial Hirsch conjecture*, which states that the graph of a polytope with m facets has diameter polynomial in m , is still open.

We conclude this section with Balinski's theorem about the connectivity of $G(\mathcal{P})$.

Theorem 10.20 (Balinski). *For any d -dimensional polytope \mathcal{P} , its graph $G(\mathcal{P})$ is d -connected.*

Proof. Let $\mathcal{P} = \text{conv}(V) \subseteq \mathbb{R}^d$ where V is the vertex set of \mathcal{P} , with $|V| \geq d + 1$. We want to show that deleting any subset $S \in \binom{V}{d-1}$ does not disconnect $G(\mathcal{P})$.

Let us fix a vertex $v_0 \in V \setminus S$ and a hyperplane $h : c^\top x = \zeta$ that goes through $S \cup \{v_0\}$. Such a plane must exist because any d points in \mathbb{R}^d is contained in some hyperplane. Let ζ_{\min} and ζ_{\max} be the minimum and maximum values that the linear function $c^\top x$ can attain on \mathcal{P} , respectively; note that $\zeta_{\min} \leq \zeta \leq \zeta_{\max}$.

Let F_{\min} and F_{\max} be the faces supported by the hyperplanes $c^\top x = \zeta_{\min}$ and $c^\top x = \zeta_{\max}$, respectively. Now consider an arbitrary vertex $v \in V \setminus S$.

- If $c^\top v \geq \zeta$, then by Proposition 10.19, either $v \in V(F_{\max})$ or there is a path from v to $V(F_{\max})$ such that the function value $c^\top x$ strictly increases in each hop. In particular, the path avoids the set S because $c^\top x = \zeta$ for all $x \in S$.
- If $c^\top v \leq \zeta$, then by a symmetric argument, either $v \in F_{\min}$ or there is a path from v to $V(F_{\min})$ that avoids S .

Moreover, since $c^\top v_0 = \zeta$ we know that v_0 connects to *both* F_{\min} and F_{\max} without going through set S .

Finally, observe that F_{\min} and F_{\max} are lower dimensional polytopes, so by induction both $G(F_{\min})$ and $G(F_{\max})$ are connected. Therefore we may conclude that all vertices in $V \setminus S$ are connected without going through S . □

10.4.2 The Face Lattice

The graph of a polytope concerns only the 0-faces (vertices) and 1-faces (edges). More generally, we can collect all the faces of a polytope \mathcal{P} and order them by inclusion. That is, $F \leq G$ if $F \subseteq G$; and $F < G$ if $F \subset G$. This partially ordered set (or poset) is called the **face lattice** of \mathcal{P} . Posets are usually drawn as *Hasse diagrams* where larger elements appear higher up. Two elements $F < G$ are linked by a line if there is no $H : F < H < G$. For example, the face lattice of the 3-dimensional cube is depicted in Figure 10.6.




Figure 10.6: *The cube (left) and its face lattice (right). Faces are named with the labels of their vertices.*

What makes this poset a *lattice* is the following property. For any two faces F and G , there is

- a unique maximal face E such that $E \leq F, G$ (called their *meet*); and
- a unique minimal face H such that $H \geq F, G$ (called their *join*).

Clearly the meet of F and G is exactly $F \cap G$ (which is indeed a face by Lemma 10.9, with vertex set $V(F) \cap V(G)$). It may be tempting to think that the join of F and G is $\text{conv}(V(F) \cup V(G))$, but that is not a face in general. The join turns out to be something more intricate. We can already observe this in the face lattice of a cube (Figure 10.6). The join of the edges 12 and 13, for example, is the face with four vertices 1234. The following exercise asks you to prove the existence of a join, implicitly.

Exercise 10.21. *In general, a poset is a pair (\mathcal{F}, \leq) . Here \leq is a partial order over \mathcal{F} , meaning that it is reflexive ($F \leq F$ always holds), antisymmetric ($F \leq G$ and $G \leq F$ implies $F = G$) and transitive ($F \leq G$ and $G \leq H$ implies $F \leq H$).*

An element $F \in \mathcal{F}$ is maximal if there is no element $G > F$. Similarly, it is minimal if there is no element $G < F$.

An element E is a maximal lower bound of F and G if $E \leq F, G$ but no element $E' > E$ has this property. If there is only one such E , then we call it the meet of F and G . Similarly, an element H is a minimal upper bound of F and G if $F, G \leq H$

but no element $H' < H$ has this property. If there is only one such H then we call it the join of F and G .

Now for the actual exercise: Let (\mathcal{F}, \leq) be a finite poset with a unique maximal element 1. Further suppose that every two elements F and G have a meet. Prove that then also every two elements F and G have a join!

For other fine-grained properties of the face lattice, see [12, Theorem 2.7].

The face lattice stores the combinatorial information of a polytope. Two polytopes are called *combinatorially equivalent* if they have isomorphic face lattices [12, Section 2.2]. Combinatorially equivalent polytopes may geometrically look different. For example, all triangles in the plane are combinatorially equivalent, but some are equilateral while others can be long and skinny.

10.4.3 Polarity

For every polytope $\mathcal{P} \ni 0$, there is a so-called **polar polytope** $\mathcal{P}^\Delta \ni 0$ whose face lattice is that of \mathcal{P} but turned upside down [12, Theorem 2.11]. This means that the vertices of \mathcal{P} correspond to facets of \mathcal{P}^Δ , edges of \mathcal{P} to ridges of \mathcal{P}^Δ , and so on.

If $\mathcal{P} = \text{conv}(\mathcal{P})$, then its polar polytope can be constructed as

$$\mathcal{P}^\Delta = \bigcap_{p \in \mathcal{P}} h_p^- \quad \text{where} \quad h_p^- := \left\{ x \in \mathbb{R}^d : \sum_{i=1}^d p_i x_i \leq 1 \right\}.$$

Geometrically, going to the polar polytope corresponds to replacing a point (part of the description of the \mathcal{V} -polytope \mathcal{P}) with a halfspace (part of the description of the \mathcal{H} -polytope \mathcal{P}^Δ). This operation is called *inversion at the unit sphere*; see Figure 10.7. It can be shown that $\mathcal{P}^{\Delta\Delta} = \mathcal{P}$.




Figure 10.7: The polar halfspace h_p^- has distance $1/\|p\|$ from the origin 0 and is perpendicular to the vector p .

We can also “polarize” \mathcal{P} even if it does not contain the origin: simply choose the center of inversion at some interior point of \mathcal{P} . Depending on which point we choose, \mathcal{P}^Δ can look different geometrically, but the combinatorial structure (i.e. face lattice) is nevertheless the same.

We have already seen some pairs of polar polytopes: In fact, each platonic solid is polar to another one (Figure 10.8). As a sanity check, the dodecahedron has 12 facets (hence its name), 30 edges and 20 vertices; its polar, the icosahedron, has 20 facets (hence its name), 30 edges and 12 vertices.




Figure 10.8: *Polarities among the platonic solids: the tetrahedron is polar to itself (first column); cube and octahedron are polar to each other (second and third columns); dodecahedron and icosahedron are polar to each other (fourth and fifth columns).*

Three of the platonic solids generalize to polytopes in arbitrary dimension d , and we have already encountered two of these in Section 10.3: simplices and hypercubes. Exercise 10.22 below asks you to show that simplices are polar to simplices. What polytopes are polar to hypercubes? These are called *cross-polytopes* which generalizes the octahedron. The standard d -dimensional cross-polytope is the convex hull of the $2d$ points $(\pm 1, 0, \dots, 0), \dots, (0, 0, \dots, \pm 1)$. Equivalently, we may represent it as the intersection of 2^d halfspaces $\{x \in \mathbb{R}^d : \sum_{i=1}^d h_i x_i \leq 1\}$ for $h \in \{-1, 1\}^d$.

Exercise 10.22. *Argue that the face lattice of a d -simplex is isomorphic to the Boolean cube $(\{0, 1\}^{d+1}, \leq)$. Conclude that d -simplex is polar to itself. (Hint: Exercise 10.18.)*

Polarity sometimes yields surprisingly short proofs that would otherwise require a non-trivial argument. Below is an example.

Lemma 10.23. *Every proper face is contained in some facet.*

Proof. Via polarity, the statement for \mathcal{P} translates to “every proper face in \mathcal{P}^Δ contains one or more vertices”. The latter follows from Lemma 10.7. \square

10.5 Simplicial and Simple Polytopes

Let us return to the important question that we asked earlier in this chapter:

How many facets can a d -dimensional polytope with n vertices have?

As we have discussed, the answer is $\Theta(n)$ for $d = 2, 3$. For general d , we have an upper bound of $O(n^d)$ from Exercise 10.8, which is an overestimate for $d = 2, 3$ already. The full answer will come later in Chapter 11, but let us make a general observation right now. To address the question, we can actually restrict our attention to **simplicial polytopes**. These are d -dimensional polytopes whose facets are all simplices (or more specifically, $(d - 1)$ -simplices). For example, the octahedron is simplicial since all its facets are triangles (2-simplices), whereas the dodecahedron is not since its facets are pentagons.

Fixing dimension d and the number n of vertices, the number of facets can only be maximized by a simplicial polytope. The reason is that a non-simplicial polytope can be “made simplicial” by slightly and randomly perturbing its vertices. Intuitively, each non-simplex facet “breaks apart” and gets replaced by several simplex facets. More formally, with probability 1, all subsets of $d + 1$ vertices become affinely independent, thus forcing every facet to contain exactly d vertices (otherwise its dimension would not be $d - 1$). Hence each facet is a $(d - 1)$ -simplex now. One can show that the original facets injectively maps to the new facets, so the number of facets cannot decrease.³

Let’s illustrate this in the cube $[-1, 1]^3$. Suppose that we push the two vertices $(-1, -1, -1)$ and $(1, 1, 1)$ “slightly inwards” so that they become $(-1 + \epsilon, -1 + \epsilon, -1 + \epsilon)$ and $(1 - \epsilon, 1 - \epsilon, 1 - \epsilon)$, respectively, for some small $\epsilon > 0$, then we obtain the simplicial polytope in Figure 10.9. Similarly, for the dodecahedron, each pentagon facet gets replaced by three triangles when we slightly perturb the vertices.




Figure 10.9: *Perturbing the cube vertices: by pushing the two diagonal vertices slightly inwards, each square facet breaks up into two triangles, and the resulting polytope is simplicial.*

³In fact, it will strictly increase. See [12, Lemma 8.24] for a formal statement and reference to a proof.

Exercise 10.24. *What happens if we move the two vertices “slightly outwards” so that they become $(-1 - \varepsilon, -1 - \varepsilon, -1 - \varepsilon)$ and $(1 + \varepsilon, 1 + \varepsilon, 1 + \varepsilon)$, respectively? Draw the resulting simplicial polytope!*

Simplicial polytopes not only maximize the number of facets; they are also very nice structurally. In fact, not only their facets but also *all* their faces are simplices!

Proposition 10.25. *A d -dimensional polytope is simplicial if and only if every k -face is a k -simplex, for $0 \leq k \leq d - 1$.*

Proof. The (\Leftarrow) direction is trivial. For the (\Rightarrow) direction, Lemma 10.23 states that every k -face F is contained in some facet F' , which is a simplex by assumption. In particular, $V(F) \subseteq V(F')$ is a set of affinely independent points. So F is a k -simplex. \square

We can also define the polar notion of simplicial polytopes. A polytope is **simple** if every vertex is incident to d edges. As polarity transform turns the face lattice upside down, a polytope is simple if and only if its polar polytope is simplicial. Checking Figure 10.8 again, the tetrahedron is both simple and simplicial, the octahedron as well as the icosahedron are simplicial, and their polars—the cube and the dodecahedron—are simple. Via polarity, an alternative way to phrase our initial question is:

How many vertices can a d -dimensional polytope with n facets have?

The count is maximized only by the simple polytopes.

Exercise 10.26. *Characterize all polytopes in \mathbb{R}^3 that are both simple and simplicial.*

10.6 High-Dimensional Delaunay Triangulations

In discussing Delaunay triangulations and proving the termination of the Lawson flip algorithm in Section 6.3, we have argued that every triangulation in the plane gives rise to a “lifted surface” that can pointwise only decrease in height under a Lawson flip, so that eventually no Lawson flip is possible any more. In this section we discuss more systematically what the lifted surface actually is after the algorithm terminates, that is, when the triangulation has become Delaunay. In fact, we want to generalize this to arbitrary dimension d .

We will give the big picture upfront, borrowing the very neat Figure 10.10 from Hang Si [9]. Let us consider a planar point set in general position (no three points on a line, no four points on a circle), so its Delaunay triangulation is unique by Corollary 6.18. Imagine lifting the points to the unit paraboloid in \mathbb{R}^3 and consider the convex hull of the lifted points—a polytope in \mathbb{R}^3 . Its lower facets (triangles by general position), when projected back to \mathbb{R}^2 , must satisfy the empty circle property (cf. Lemma 6.12) and thus yield the Delaunay triangulation. See the lower right part in Figure 10.10.

So the “lifted surface” when the Lawson flip algorithm terminates is exactly the *lower convex hull* of the lifted points. This also means that we can reduce the computation




Figure 10.10: *Triangulations in \mathbb{R}^d as projections of polytopes in \mathbb{R}^{d+1}*

of the Delaunay triangulation to computing a convex hull in \mathbb{R}^3 . We will formally state and prove this relation for general dimension soon.

Figure 10.10 shows more. In the upper right part, we see what happens when we project the *upper* facets back to \mathbb{R}^2 . The result is called the *farthest-point Delaunay triangulation*. It is generally not a triangulation of the point set, but only of the extreme points. Each triangle in this triangulation is “anti-Delaunay” in the sense that its circumcircle contains *all* other points; see Exercise 10.41 below. The left part of Figure 10.10 shows what happens if we lift the points not onto the unit paraboloid but in some arbitrary way. The convex hull of the lifted points is still a polytope, and if it is simplicial, we can recover two triangulations in the plane by projecting the lower and upper facets back to \mathbb{R}^2 , respectively. Such triangulations are called *regular*; the (farthest-point) Delaunay triangulation is a specific regular triangulation.

After the pictorial outline, we will now formalize the intuition. In Definition 6.1, we have introduced triangulations of point sets in the plane. We can generalize it to higher dimensions in a straightforward way, replacing “triangles” by “d-simplices”. We still call it a triangulation, for lack of a better name derived from the word “simplex”.

Definition 10.27. *A triangulation of a finite point set $P \subset \mathbb{R}^d$ is a collection \mathcal{T} of d-simplices, such that*

$$(1) \text{conv}(P) = \bigcup_{T \in \mathcal{T}} T;$$

(2) $P = \bigcup_{T \in \mathcal{T}} V(T)$; and

(3) for every distinct pair $T, T' \in \mathcal{T}$, the intersection $T \cap T'$ is a face of both.⁴

For $d = 2$ we recover Definition 6.1. Also for $d = 1$ the definition makes sense, as a point set $\{a_1, a_2, \dots, a_n\}$ in \mathbb{R}^1 (assuming $a_1 < a_2 < \dots < a_n$) has a unique triangulation $\mathcal{T} = \{[a_i, a_{i+1}] : 1 \leq i < n\}$.

However, at the moment it is not clear whether every point set in \mathbb{R}^d has a triangulation, for $d \geq 3$. Anyway, we go ahead and define Delaunay triangulations in the same way as before.

Definition 10.28. A triangulation \mathcal{T} of a finite point set $P \subset \mathbb{R}^d$ is a Delaunay triangulation, if the circumsphere of every d -simplex $T \in \mathcal{T}$ is empty of points from P .

What is the circumsphere of a d -simplex? This is the unique sphere that contains all its $d + 1$ vertices. Before you can even question about its existence and uniqueness, let us prove it.

Lemma 10.29. Let $S \subset \mathbb{R}^d$ be a set of $d + 1$ affinely independent points. Then there exists a unique sphere containing S .

Proof. Recall that a sphere with center $c \in \mathbb{R}^d$ and radius $r \geq 0$ is formally defined as the set $\{x \in \mathbb{R}^d : \|x - c\| = r\}$. Squaring the condition, we are looking for a (unique) point $c \in \mathbb{R}^d$ and a (unique) number $r \geq 0$ such that

$$\|q - c\|^2 = r^2, \quad \forall q \in S \quad (10.30)$$

As usual, we understand a point $x \in \mathbb{R}^d$ as a column vector. Then $x^\top y$ is the scalar product $\sum_{i=1}^d x_i y_i$ of two points $x, y \in \mathbb{R}^d$. With this, the previous system of equations is equivalent to

$$\|q\|^2 = 2q^\top c + \underbrace{r^2 - \|c\|^2}_{=: \alpha}, \quad \forall q \in S. \quad (10.31)$$

In still other words,

$$\|q\|^2 = (q^\top, 1) \begin{pmatrix} 2c \\ \alpha \end{pmatrix}, \quad \forall q \in S.$$

Stacking the $d + 1$ equations row by row, this is a linear system of the form $b = B \begin{pmatrix} 2c \\ \alpha \end{pmatrix}$ where $b \in \mathbb{R}^{d+1}$ and $B \in \mathbb{R}^{(d+1) \times (d+1)}$, one row for each $q \in S$. As the points $q \in S$ are affinely independent, the rows of B are linearly independent (Proposition 5.4) and so B is invertible. So there is a unique $c \in \mathbb{R}^d$ and a unique $\alpha \in \mathbb{R}$ satisfying (10.31), which uniquely determine $r^2 := \alpha + \|c\|^2$ and satisfy (10.30). Note that such r^2 must be positive because the left hand side of this satisfied equation (10.30) is always positive. \square

⁴ Note that this also allows for $T \cap T' = \emptyset$, since \emptyset is a face of every polytope.

Next we want to show that there is a unique Delaunay triangulation, assuming sufficiently general position. This means that no $d + 1$ points lie on a common hyperplane, and no $d + 2$ points lie on a common sphere.

As a preparation, we first define the concept of a Delaunay simplex.

Definition 10.32. *Let $P \subset \mathbb{R}^d$ be a set of points in general position. A simplex $\text{conv}(S)$ where $S \in \binom{P}{d+1}$ is called a **Delaunay simplex** for P if the circumsphere of S is empty of points from P .*

Next comes the crucial insight. Generalizing Section 6.3, we show that Delaunay simplices correspond to “lower” facets of a polytope in one dimension higher, namely the convex hull of the lifted points. For $p = (p_1, \dots, p_d) \in \mathbb{R}^d$, we define the *lifted point*

$$\ell(p) := (p, \|p\|^2) = (p_1, \dots, p_d, \|p\|^2) \in \mathbb{R}^{d+1}. \tag{10.33}$$

For $d = 2$, this is the standard lifting map that raises points in the plane to the unit paraboloid in \mathbb{R}^3 . The following lemma naturally extends Lemma 6.12.

Lemma 10.34. *For any sphere $S \subset \mathbb{R}^d$, there is an upward hyperplane⁵ $h \subset \mathbb{R}^{d+1}$ such that the following property holds. A point $q \in \mathbb{R}^d$ is on/outside/inside S if and only if the lifted point $\ell(q) \in \mathbb{R}^{d+1}$ is on/above/below h .*

Conversely, for any upward hyperplane $h \subset \mathbb{R}^{d+1}$ that intersects the unit paraboloid, there is a sphere $S \subset \mathbb{R}^d$ such that the aforementioned property holds.

Proof. We have already done most of the work in the proof of Lemma 10.29. Given a sphere $S \subset \mathbb{R}^d$ with center c and radius r , let us denote $\alpha := r^2 - \|c\|^2$. Along the same lines of deriving (10.31), we have

$$\begin{aligned} \|q\|^2 = 2q^\top c + \alpha &\iff q \text{ on } S, \\ \|q\|^2 > 2q^\top c + \alpha &\iff q \text{ outside } S, \\ \|q\|^2 < 2q^\top c + \alpha &\iff q \text{ inside } S. \end{aligned} \tag{10.35}$$

Recall that $\ell(q) = (q, \|q\|^2)$, so the formulas may be rephrased as

$$\begin{aligned} (-2c^\top, 1)\ell(q) = \alpha &\iff q \text{ on } S, \\ (-2c^\top, 1)\ell(q) > \alpha &\iff q \text{ outside } S, \\ (-2c^\top, 1)\ell(q) < \alpha &\iff q \text{ inside } S. \end{aligned} \tag{10.36}$$

In other words, the lifted point $\ell(q)$ is on/above/below the upward hyperplane $h := \{x \in \mathbb{R}^{d+1} : (-2c^\top, 1)x = \alpha\}$ in the respective cases.

Conversely, given an upward hyperplane $h \subset \mathbb{R}^{d+1}$, let $h_1, \dots, h_d, 1, h_{d+2}$ be its parameters. Define $c_i := -h_i/2$, $\alpha := h_{d+2}$ and $r^2 := \alpha + \|c\|^2$, and consider the *formal* sphere $S := \{x \in \mathbb{R}^d : \|x - c\| = r\}$. It is not yet clear that $r^2 \geq 0$, or $S \neq \emptyset$, but the

⁵ Being upward means that the coefficient for the last coordinate is positive. By scaling the equation appropriately, we may assume that the coefficient is exactly one. Geometrically, the normal vector of such hyperplane is pointing upward.

derivations (10.35) and (10.36) hold anyway, so the desired property definitely hold. We just need to show $\mathcal{S} \neq \emptyset$. To this end, recall from assumption that h intersects the unit paraboloid, thus the “=” case in (10.36) does happen, which certifies that there is *some* point on \mathcal{S} . \square

Corollary 10.37. *Let $P \subset \mathbb{R}^d$ be a finite point set. Denote by $\ell(P) = \{\ell(p) : p \in P\}$ the set of lifted points. Then the polytope $\mathcal{P} := \text{conv}(\ell(P))$ has vertex set $\ell(P)$.*

Proof. By definition $V(\mathcal{P}) \subseteq \ell(P)$, so it remains to show that $\ell(p)$ is vertex for all $p \in P$. To this end, apply Lemma 10.34 to the singleton $\mathcal{S} = \{\ell(p)\}$ (a sphere with center $\ell(p)$ and radius 0!) and get a hyperplane h . Every point $q \in P \setminus p$ is outside the sphere \mathcal{S} , so its lifting $\ell(q)$ is above the hyperplane h . Hence h supports $\{\ell(p)\}$ by Lemma 10.4. \square

Lemma 10.38. *Let $P \subset \mathbb{R}^d$ be a finite point set in general position. Then $\mathcal{P} = \text{conv}(\ell(P))$ is a simplicial polytope in \mathbb{R}^{d+1} . Moreover, for any subset $S \subseteq P$, the following two statements are equivalent.*

- $\text{conv}(S)$ is a Delaunay simplex for P .
- $\text{conv}(\ell(S))$ is a lower facet of \mathcal{P} , meaning that it is a facet supported by some upward hyperplane.

Proof. That \mathcal{P} is simplicial follows from general position. Indeed, every facet of \mathcal{P} is a d -face, so it contains at least $d+1$ (affinely independent) vertices. But it cannot contain more: All the vertices, necessarily in the form $\ell(p)$, $p \in P$ by Corollary 10.37, are lying on a common hyperplane in \mathbb{R}^{d+1} , so their projections onto \mathbb{R}^d are on a common sphere by Lemma 10.34. General position requires the number to be less than $d+2$.

Now we proceed to the “moreover” part. Let $\text{conv}(S)$ be a Delaunay simplex, so S consists of $d+1$ affinely independent points whose circumsphere is empty of points from P . Applying Lemma 10.34 on this sphere, there is an upward hyperplane h such that $\ell(P) \cap h = \ell(S)$ and $\ell(P) \subset h^+$. So h supports $\text{conv}(\ell(S))$ by Lemma 10.4. Note that $\ell(S)$ consists of $d+1$ affinely independent points, so the face $\text{conv}(\ell(S))$ has dimension d and is a (lower) facet, indeed.

Conversely, assume that $\text{conv}(\ell(S))$ is a lower facet (a d -simplex since the polytope is simplicial), supported by some upward hyperplane h . This time we apply Lemma 10.34 on h , and obtain a sphere that goes through S and satisfies the empty property. Finally, S is indeed a simplex by general position (no $d+1 = |S|$ points on a common hyperplane). \square

From this correspondence, we may obtain the existence of a unique Delaunay triangulation for a finite point set $P \subset \mathbb{R}^d$ in general position.

Theorem 10.39. *Let $P \subset \mathbb{R}^d$ be a finite point set in general position. Then the collection \mathcal{T} of all Delaunay simplices for P is the unique Delaunay triangulation of P .*

Proof. Suppose \mathcal{T} is a triangulation. Then it would be a Delaunay triangulation by definition. The uniqueness also follows: Any other Delaunay triangulation consists of Delaunay simplices, thus a proper subset of \mathcal{T} ; but then it cannot cover $\text{conv}(P)$ in full, a contradiction.

It remains to prove that \mathcal{T} is a triangulation, so let's look at the three properties in Definition 10.27. Denote by $\mathcal{P} = \text{conv}(\ell(P))$ the convex hull of the lifted points (a polytope in \mathbb{R}^{d+1}).

(1) $\text{conv}(P) = \bigcup_{T \in \mathcal{T}} T$. Let $q \in \text{conv}(P)$, and our goal is to find a simplex $T \in \mathcal{T}$ that contains q . Consider a vertical⁶ line in \mathbb{R}^{d+1} from $(q, -\infty)$ to (q, ∞) . Since $q \in \text{conv}(P)$, the line must intersect $\text{conv}(\ell(P)) = \mathcal{P}$ in a non-empty closed interval (it is an interval since \mathcal{P} is convex and compact). So let us choose the minimum height $t \in \mathbb{R}$ such that $(q, t) \in \mathcal{P}$. Then (q, t) is on the boundary of \mathcal{P} and hence contained in one or more facets by Corollary 10.16; one of these must be a lower facet $\text{conv}(\ell(S))$ by Exercise 10.40 below. So $q \in \text{conv}(S)$. But we know $\text{conv}(S) \in \mathcal{T}$ by Lemma 10.38, and this is the simplex we are looking for.

We have shown $\text{conv}(P) \subseteq \bigcup_{T \in \mathcal{T}} \text{conv}(T)$. The reverse inclusion follows from $\text{conv}(S) \subseteq \text{conv}(P)$ for all $S \in \binom{P}{d+1}$.

(2) $\mathcal{P} = \bigcup_{T \in \mathcal{T}} V(T)$. The inclusion \supseteq is trivial. Now for the other inclusion, let $p \in \mathcal{P}$. We claim that $\ell(p)$ is the vertex of some lower facet of \mathcal{P} . Via Lemma 10.38 this implies that p is the vertex of some $T \in \mathcal{T}$.

Here is the key observation: $\min\{t \in \mathbb{R} : (p, t) \in \mathcal{P}\} = \|p\|^2$. Indeed, for $t = \|p\|^2$ we have $(p, t) = \ell(p) \in \mathcal{P}$. But if $t < \|p\|^2$ then (p, t) is outside the convex “bowl” $\mathcal{U} = \{x \in \mathbb{R}^{d+1} : x_{d+1} \geq \sum_{i=1}^d x_i^2\}$, whereas $\mathcal{P} = \text{conv}(\ell(P)) \subset \mathcal{U}$.

By the argument for (1), vertex $\ell(p)$ is hence contained in some lower facet $\text{conv}(\ell(S))$ of \mathcal{P} and is then also a vertex of this facet. So the claim is proved.


(3) **The intersection of any two simplices $T, T' \in \mathcal{T}$ is a face of both.** This follows from the general structure of polytopes. Let F and F' be the lower facets of \mathcal{P} corresponding to the Delaunay simplices T and T' . By Lemma 10.9, the intersection $F \cap F'$ is a face of \mathcal{P} with vertex set $\ell(V(T)) \cap \ell(V(T')) = \ell(V(T) \cap V(T'))$, hence $F \cap F' = \text{conv}(\ell(V(T) \cap V(T')))$. Projecting back onto \mathbb{R}^d we see $T \cap T' = \text{conv}(V(T) \cap V(T'))$. This is a face of both simplices T and T' , since every subset of vertices of a simplex defines a face (Exercise 10.18). □

Exercise 10.40. Let $\mathcal{P} \subset \mathbb{R}^{d+1}$ be a polytope and $(q, t) \in \mathbb{R}^{d+1}$ such that $(q, t) \in \mathcal{P}$ but $(q, t') \notin \mathcal{P}$ for $t' < t$. Prove that (q, t) is contained in some lower facet of \mathcal{P} .

Exercise 10.41. Let $P \subset \mathbb{R}^d$ be a finite set of points in convex position (every point is extreme), and in general position (no $d + 1$ points on a hyperplane, no $d + 2$ on a

⁶ In high dimension, the word “vertical” should read “along the last axis”. Similarly, the word “height” should read “the last coordinate”.

sphere). A farthest-point Delaunay triangulation of P is a triangulation \mathcal{T} of P with the property that the circumsphere of every d -simplex T in \mathcal{T} contains all points $P \setminus V(T)$:



Prove that P has a unique farthest-point Delaunay triangulation; Figure 10.10 provides the intuition. The name comes from the fact that in the plane, the farthest-point Delaunay triangulation is dual to the farthest-point Voronoi diagram, the subdivision of the plane into regions with the same farthest point.

10.7 Complexity of 4-Dimensional Polytopes

The **complexity** of a polytope is defined as the number of faces. Indeed, if we talk about computing a polytope, we typically mean that we want to compute its face lattice. In dimensions $d = 2, 3$, each polytope with n vertices has complexity $\Theta(n)$. We have also seen that for $d = 4$, the complexity is bounded by $O(n^4)$ (Exercise 10.8). But can we actually have superlinear complexity for $d = 4$, or does the linear behavior in dimensions $d = 2, 3$ continue?

Using the previously derived connection to 3-dimensional Delaunay triangulations, we can answer this question.

Theorem 10.42. *For every even natural number $n \geq 4$, there exists a 4-dimensional simplicial polytope with n vertices and at least $\binom{n}{2} - 1 = \Theta(n^2)$ facets.*

Moreover, this polytope also has $\Theta(n^2)$ edges which is asymptotically maximal since Exercise 10.8 implies that there are $O(n^2)$ edges. In particular, vertex-edge graphs of 4-dimensional polytopes can be dense and highly non-planar. They can even be complete [12, Corollary 0.8]. This may be somewhat counter-intuitive, as it seems to require the edges “cutting through” the polytope which they obviously cannot. On the other hand, 4 dimensions are counterintuitive *per se*, so let’s not worry too much about intuition here.

Proof. We construct a point set $P \subset \mathbb{R}^3$ in general position, $|P| = n$, for which there are at least $\binom{n}{2} - 1$ Delaunay simplices. By Lemma 10.38, the convex hull of the lifted point set $\ell(P)$ is a 4-dimensional simplicial polytope with at least $\binom{n}{2} - 1$ (lower) facets.

Let ℓ_1, ℓ_2 be two skew (non-parallel, non-intersecting) lines in \mathbb{R}^3 . We choose a set P_1 of $n/2$ points on ℓ_1 , and another set P_2 of $n/2$ points on ℓ_2 . Then we set $P = P_1 \cup P_2$, after slightly perturbing all points to ensure general position.

The claim is that for any points $p, q \in P_1$ consecutive along ℓ_1 , and points $r, s \in P_2$ consecutive along ℓ_2 , their convex hull $\text{conv}\{p, q, r, s\}$ is a Delaunay simplex. (See the cartoonish Figure 10.11.) As there are $\binom{n}{2}$ ways to choose such p, q, r, s , there are at least this many Delaunay simplices.




Figure 10.11: Proof of Theorem 10.42

It remains to prove the claim. The points p, q, r, s are affinely independent (by general position) and hence have a unique circumsphere. The line ℓ_1 intersects this sphere in exactly the line segment \overline{pq} ; but p, q are consecutive along ℓ_1 , so no point of P_1 lies inside the sphere. For the same reason, no point of P_2 lies inside. As the sphere is empty, $\text{conv}\{p, q, r, s\}$ is a Delaunay simplex. \square

It is actually the case that a 4-dimensional polytope with n vertices has $O(n^2)$ facets, so the lower bound provided by Theorem 10.42 is asymptotically best possible. We will postpone the full account to the later Chapter 11, where we give a tight upper bound on the number of facets that a d -dimensional polytope with n vertices can have.

10.8 High Dimensional Voronoi Diagrams

Using lifting map, we can also relate the Voronoi diagram of a finite point set $P \subset \mathbb{R}^d$ with the facets of some *polyhedron* in \mathbb{R}^{d+1} . In fact this is what Theorem 8.16 did for $d = 2$, without explicitly mentioning polyhedra. Here we simply reprove this theorem for general d . No new idea appear, so the reader is invited to consider it as a repetition of Section 8.4, but formulated in the more abstract language of polyhedra and replacing “2” by “ d ”.

Let us start by generalizing Voronoi cells to higher-dimensions which is a straightforward adaptation of Definition 8.3.

Definition 10.43. Let $P \subset \mathbb{R}^d$ be a finite point set. The Voronoi cell of point $p \in P$ is defined as

$$V_P(p) := \{q \in \mathbb{R}^d : \|q - p\| \leq \|q - p'\| \text{ for all } p' \in P\}.$$

In words, $V_P(p)$ is the set of points in \mathbb{R}^d for which p is a (not necessarily unique) closest point among all points in P .

Theorem 10.44. *Let $P \subset \mathbb{R}^d$ be a finite point set. For each $p \in P$, we define a hyperplane*

$$h_p = \left\{ x \in \mathbb{R}^{d+1} : x_{d+1} - \sum_{i=1}^d 2p_i x_i = -\|p\|^2 \right\}.$$

Consider the polyhedron $\mathcal{P} := \bigcap_{p \in P} h_p^+$ in \mathbb{R}^{d+1} . Then every h_p supports a lower facet of \mathcal{P} . Moreover, for all $q \in \mathbb{R}^d$, the following two statements are equivalent.

- (i) $q \in V_P(p)$.
- (ii) $(q, t) \in h_p$, where $t \in \mathbb{R}$ is the minimum value such that $(q, t) \in \mathcal{P}$.

Pictorially, condition (ii) means that the vertical ray emanating up from $(q, -\infty)$ hits the polytope at the lower facet $\mathcal{P} \cap h_p$. Hence the theorem says that the Voronoi cell $V_P(p)$ is simply the vertical projection of the facet $\mathcal{P} \cap h_p$ back to \mathbb{R}^d . If we project all the facets of \mathcal{P} to \mathbb{R}^d , we obtain the Voronoi diagram of P . Figure 10.12, borrowed from the book by Joswig and Theobald [7, Page 87], visualizes this for $d = 3$.




Figure 10.12: *A part of the polyhedron $\mathcal{P} \subset \mathbb{R}^3$ in Theorem 10.44, and the Voronoi diagram as the projections of its facets to \mathbb{R}^2 .*

Proof. We first show that all h_p are actually facet-supporting hyperplanes. For this, it suffices to show that none of the halfspaces h_p^+ is redundant; see the remark above Corollary 10.16.

The mysterious-looking hyperplanes h_p are actually our old friends! They appeared implicitly in the proof of Corollary 10.37. The hyperplane h_p is obtained exactly by applying Lemma 10.34 on the singleton sphere $\{\ell(p)\}$, where ℓ denotes the lifting map. Hence $\ell(p)$ is on h_p and strictly above other h_q . So it is in the interior of the polyhedron $\bigcap_{q \in P \setminus \{p\}} h_q^+$ but on the boundary of $\bigcap_{q \in P} h_q^+$, thus the hyperplane h_p is not redundant.

Next we claim that the vertical distance from $\ell(q)$ to h_p is precisely $\|q - p\|^2$ (cf. Lemma 8.14 and Figure 8.7). Indeed, $\ell(q) = (q, \|q\|^2)$ has height $\|q\|^2$, and its vertical projection onto h_p has height

$$x_{d+1} = \sum_{i=1}^d 2p_i q_i - \|p\|^2 = 2p^\top q - \|p\|^2.$$

So the vertical distance is $\|q\|^2 - 2p^\top q + \|p\|^2 = \|q - p\|^2$.

Now we can show the equivalence of (i) and (ii). Given any point $q \in \mathbb{R}^d$, $\ell(q)$ is on or above the hyperplanes h_p . So by the claim we have the following chain of equivalence:

- $q \in V_P(p)$.
- The vertically closest hyperplane to $\ell(q)$ is h_p .
- Projecting $\ell(q)$ vertically onto the hyperplanes, the highest point is on h_p .
- Raising $(q, -\infty)$ vertically until we hit \mathcal{P} , the hitting point is on h_p .
- $(q, t) \in h_p$. □

10.9 Regular Triangulations and the Secondary Polytope

We have seen in Section 10.6 that the Delaunay triangulation of a set P of points in \mathbb{R}^d can be constructed as the projection of the lower convex hull of the polytope in \mathbb{R}^{d+1} defined by lifting the points in P to the (generalized) unit paraboloid. As already hinted at in Section 10.6, we can get other (partial) triangulations through other liftings. In order to make this precise, consider a set $P = \{p_1, \dots, p_n\}$ of points in \mathbb{R}^d and let $\omega : P \rightarrow \mathbb{R}$ be a *height function*. Denote by $P^\omega := \{(p_1, \omega(p_1)), \dots, (p_n, \omega(p_n))\}$ the *lifted point set*. Consider now the polytope $\mathcal{P} = \text{conv}(P^\omega)$ and let $Q^\omega \subseteq P^\omega$ be the vertices of \mathcal{P} that lie on a lower face of \mathcal{P} . Projecting the lower convex hull of \mathcal{P} back down to \mathbb{R}^d then gives a convex subdivision on Q and thus a partial triangulation of P . Note that if Q^ω is in general position then the lower convex hull is simplicial and we get a triangulation of Q . If additionally we have that $Q^\omega = P^\omega$, that is, all vertices of \mathcal{P} lie on the lower convex hull (as is the case when lifting to the generalized unit paraboloid), then we get a triangulation of P .

Definition 10.45. A (partial) triangulation of a set P of points in \mathbb{R}^d is regular if it can be obtained by projecting the lower convex hull of a lifting of P to \mathbb{R}^{d+1} .

From Section 10.6 we can deduce that every point set admits a regular triangulation, for example the Delaunay triangulation. On the other hand, we may ask whether every triangulation is regular? While this is true for point sets in \mathbb{R}^1 (see Exercise 10.47), it fails already in the plane.

Theorem 10.46. There is a triangulation of six points in \mathbb{R}^2 that is not regular.

Proof. Consider the following point set, which is also called the *mother of all examples* (see Figure 10.13). It consists of six points that lie on two nested equilateral triangles p_1, p_2, p_3 and p_4, p_5, p_6 , where the lines p_1p_4 , p_2p_5 and p_3p_6 all intersect in a single point. Consider the triangulation \mathcal{T} obtained by drawing the edges of these two triangles, together with the edges p_1p_5 , p_2p_6 and p_3p_4 . We claim that this triangulation is not regular.

Assume for the sake of contradiction that \mathcal{T} is regular, that is, there is a lifting map ω such that \mathcal{T} is the lower convex hull of the polytope defined by the lifted points. Without loss of generality we may assume that $\omega(p_4) = \omega(p_5) = \omega(p_6) = 0$. Now, in order to get the edge p_1p_5 we need $\omega(p_2) > \omega(p_1)$. Similarly, from the edges p_2p_6 and p_3p_4 we get $\omega(p_3) > \omega(p_2)$ and $\omega(p_1) > \omega(p_3)$. Combining these three inequalities, we get $\omega(p_1) > \omega(p_1)$, which is a contradiction. \square




Figure 10.13: A non-regular triangulation of the mother of all examples.

Exercise 10.47. Show that every partial triangulation of a set P of points in \mathbb{R}^1 is regular.

While not every triangulation is regular, the set of all regular triangulations of a point set have a very nice structure: they again form a polytope, as discovered by Gelfand, Karpanov and Zelevinsky in 1989 [5, 6].

Definition 10.48. Let P be a set of points in \mathbb{R}^d and let \mathcal{T} be a triangulation of P . The GKZ-vector $\Phi_P(\mathcal{T})$ is defined as

$$\Phi_P(\mathcal{T}) := \sum_{p \in P} \sum_{F \in \mathcal{T}: p \in F} \text{vol}(F) e_p \in \mathbb{R}^{|P|},$$

where $\text{vol}()$ denotes the volume and e_i is the standard unit vector in dimension i . The polytope

$$\Sigma(P) := \text{conv}\{\Phi_P(\mathcal{T}) \mid \mathcal{T} \text{ is a triangulation of } P\} \subset \mathbb{R}^{|P|}$$

is the secondary polytope of P .

While the secondary polytope is defined as a subset of $\mathbb{R}^{|P|}$, its dimension is smaller

Theorem 10.49. Let P be a set of n affinely independent points in \mathbb{R}^d . Then the dimension of $\Sigma(P)$ is $n - d - 1$.

By definition, the secondary polytope contains all triangulations of P , but a priori they might all lie in the interior. As it turns out, it is exactly the regular triangulations that do not lie in the relative interior of any face:

Theorem 10.50 ([6]). The vertices of the secondary polytope $\Sigma(P)$ are in one-to-one correspondence with the regular triangulations of P . In other words, the GKZ-coordinates $\Phi_P(\mathcal{T})$ of a triangulation \mathcal{T} are extremal in the set of all GKZ-coordinates of triangulations of P if and only if \mathcal{T} is regular.

Now that we understand the vertices of the secondary polytope, we might ask ourselves whether the edges also have a nice interpretation. Indeed they correspond exactly to *flips*. We have seen flips of triangulations in the plane in Chapter 6, where any four points in convex position defined a flip that replaced two triangles with two other triangles, flipping between the two possible triangulations of the four points. In higher dimension the situation is analogous, but slightly more intricate.

Lemma 10.51. Let $P = (p_1, \dots, p_{d+2})$ be a set of $d + 2$ points in convex position in \mathbb{R}^d . Then P has exactly two triangulations, both of which are regular.

Proof. By Radon's theorem, P can be partitioned in a unique way into two parts P^+ and P^- such that $\text{conv}P^+ \cap \text{conv}(P^-) \neq \emptyset$. Note that as P is in convex position, each part contains at least two points. Consider now the two sets of simplices

$$\mathcal{T}^+ := \{F \subset P \mid P^+ \not\subset F\} \text{ and } \mathcal{T}^- := \{F \subset P \mid P^- \not\subset F\}.$$

As $\text{conv}P^+ \cap \text{conv}(P^-) \neq \emptyset$, any triangulation of P is contained either in \mathcal{T}^+ or in \mathcal{T}^- , but not in both. It thus remains to show that both \mathcal{T}^+ and \mathcal{T}^- are regular triangulations.

Recall that by Radon's theorem there are nonzero coefficients λ_i such that

$$\sum_{p_i \in P^+} \lambda_i p_i = \sum_{p_i \in P^-} \lambda_i p_i,$$

where $\lambda_i > 0$ for $p_i \in P^+$ and $\lambda_i < 0$ for $p_i \in P^-$. Consider now some height function $\omega : P \rightarrow \mathbb{R}$ and let

$$c = \sum_{p_i \in P} \lambda_i \omega(p_i).$$

If $c = 0$ then the lifted points lie on a hyperplane. If $c > 0$ then each point in P^+ is above the unique hyperplane spanned by the other $d + 1$ lifted points, which implies that the set of lower facets is

$$\{P \setminus \{p_i\} \mid p_i \in P^+\} = \{F \subset P \mid P^+ \not\subseteq F\},$$

which shows that \mathcal{T}^+ is regular. The analogous argument for $c < 0$ shows that \mathcal{T}^- is regular. \square

Lemma 10.51 now allows us to define a flip in general dimension analogously to the planar case: a flip on a triangulation \mathcal{T} of a set of points P in \mathbb{R}^d takes a triangulated convex subset P' of P of size $d + 2$ and replaces the triangulation on P' with the unique other one. With this we can again define the *flip graph* of triangulations of a point set.

Theorem 10.52. *Let P be a set of points in \mathbb{R}^d . Then the edge graph of the secondary polytope $\Sigma(P)$ is contained in the flip graph of regular triangulations of P .*

Using Balinski's theorem, we can thus conclude the following

Corollary 10.53. *Let P be a set of n affinely independent points in \mathbb{R}^d . Then the flip graph of regular triangulations of P is $(n - d - 1)$ -connected.*

Note that we have restricted our attention to the flip graph of *regular* triangulations. If we also include non-regular triangulations, the situation becomes much more difficult. Even in the plane, it was only shown in 2020 by Wagner and Welzl [11] that the flip graph of *all* triangulations of a set of n points in \mathbb{R}^2 is $(n/2 - 2)$ -connected (compared to the $(n - 3)$ -connectedness for regular triangulations). In dimensions 5 and above there are point sets where the flip graph of all triangulations is not connected. In dimensions 3 and 4 it is still an open problem whether the flip graph of all triangulations is connected, even for convex point sets.

Questions

51. *What is a polytope?* Give a definition and provide a few examples.
52. *What is a face of a polytope? What is a vertex, an edge, a ridge, a facet?* Give precise definitions!
53. *Can you characterize vertex-edge graphs of 3-dimensional polytopes?* Explain Steinitz' Theorem.
54. *What is a hypercube? What is a simplex?* Define these polytope and explain what their faces are.
55. *How many k-faces can a d-dimensional polytope with n vertices have?* Prove a nontrivial upper bound.
56. *What is the face lattice of a polytope?* Give a precise definition, explain what the lattice property is, and why it holds for the face lattice of a polytope.
57. *What is the polar of a given polytope?* Explain the polarity transform and how face lattices of the original polytope and its polar relate to each other. Show a pair of mutually polar polytopes and interpret the aforementioned relation in the example.
58. *What are simple and simplicial polytopes?* Explain why they are relevant with respect to counting the maximal number of facets (or vertices) that a d-dimensional polytope with n vertices (or facets) can have.
59. *How connected is the graph of a polytope?* State and prove Balinski's theorem.
60. *What is a d-dimensional (Delaunay) triangulation?* Give a precise definition.
61. *Does every point set $P \subseteq \mathbb{R}^d$ have a Delaunay triangulation?* Explain why the answer is yes under general position, why the Delaunay triangulation is unique in this case, and how you can obtain it from a polytope in one dimension higher.
62. *How many facets can a 4-dimensional polytope with n vertices have?* Prove a lower bound of $\Omega(n^2)$.
63. *What is a d-dimensional Voronoi diagram?* Give a definition and explain how the Voronoi diagram relates to a polyhedron in one dimension higher!
64. *What is a regular triangulation?* Give the definition and show that not all triangulations are regular.
65. *What is the secondary polytope?* Describe the construction and explain its connection to the flip graph of regular triangulations.

References

- [1] Kjell André and DTR, Dodecahedron.svg. <https://commons.wikimedia.org/w/index.php?curid=2231561>, licensed under CC BY-SA 3.0.
- [2] Kjell André and DTR, Hexahedron.svg. <https://commons.wikimedia.org/w/index.php?curid=2231470>, licensed under CC BY-SA 3.0.
- [3] Kjell André and DTR, Tetrahedron.svg. <https://commons.wikimedia.org/w/index.php?curid=2231463>, licensed under CC BY-SA 3.0.
- [4] Cyp and DTR, Icosahedron.svg. <https://commons.wikimedia.org/w/index.php?curid=2231553>, licensed under CC BY-SA 3.0.
- [5] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky, Newton polyhedra of principal a -determinants. *Soviet Math. Dokl.*, 40, (1990), 278–281.
- [6] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky, Discriminants of polynomials in several variables and triangulations of Newton polytopes. *Leningrad Math. J.*, 2, (1991), 449–505.
- [7] Michael Joswig and Thorsten Theobald, *Polyhedral and Algebraic Methods in Computational Geometry*, Universitext, Springer, 2013.
- [8] Francisco Santos, [A counterexample to the Hirsch Conjecture](#). *Annals of Math.*, 176/1, (2012), 383–412.
- [9] Hang Si, [On Monotone Sequences of Directed Flips, Triangulations of Polyhedra, and Structural Properties of a Directed Flip Graph](#). *CoRR*, abs/1809.09701, (2018), 1–40.
- [10] Stannered, Octahedron.svg. <https://commons.wikimedia.org/w/index.php?curid=1742116>, licensed under CC BY-SA 3.0.
- [11] Uli Wagner and Emo Welzl, Connectivity of triangulation flip graphs in the plane. *Discrete & Computational Geometry*, 68/4, (2022), 1227–1284.
- [12] Günter M. Ziegler, *Lectures on Polytopes*, vol. 152 of *Graduate Texts in Mathematics*, Springer, Heidelberg, 1994.