Chapter 9

Convex Polytopes

Recall that we have defined a *convex polytope* to be the convex hull of a finite point set $P \subset \mathbb{R}^d$ (Definition 5.7). In this chapter, we take a closer at convex polytopes and their structure; in particular, we discuss their connections to Delaunay triangulations and Voronoi diagrams. On the way, we are borrowing a lot of material from Ziegler's classical book *Lectures on Polytopes* [4]. In the sequel, we will omit the attribute *convex* whenever we talk about polytopes.

We are already somewhat familiar with polytopes in dimensions $d = 2, 3$. For $d = 2$, a polytope is just a convex polygon, see Figure 9.1.

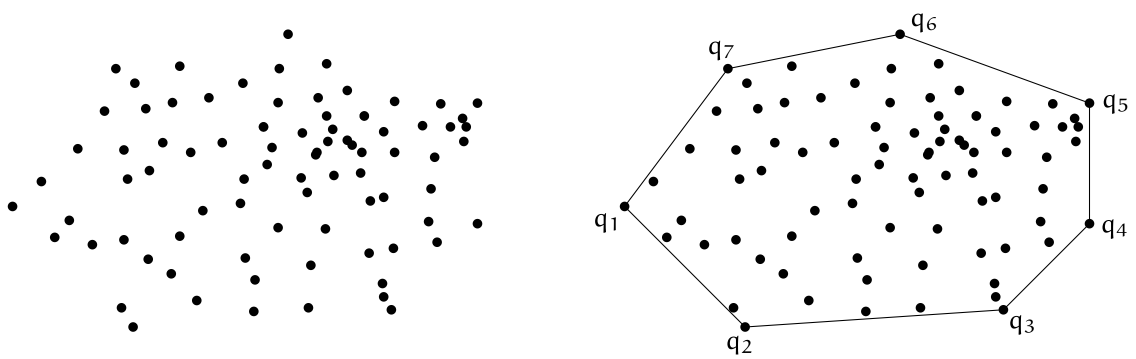

(a) Finite point set $P \subset \mathbb{R}^2$.

(b) $\mathrm{conv}(P)$, a convex polygon with vertex set $Q \subset P$

Figure 9.1: *Convex polytopes in* $\mathbb{R}^2$ *are convex polygons.*

Convex polygons are boring in the sense that they all look the same, combinatorially: the vertex-edge graph of a convex polygon with n vertices is just a simple cycle of length n. Therefore, from a graph-theoretical viewpoint, all convex polygons with n vertices are isomorphic.

The situation is more interesting for $d = 3$ where infinitely many combinatorially different polytopes exist. The most popular examples are the five platonic solids. Two of them are the *octahedron*[1] and the *dodecahedron*[2], see Figure 9.2.

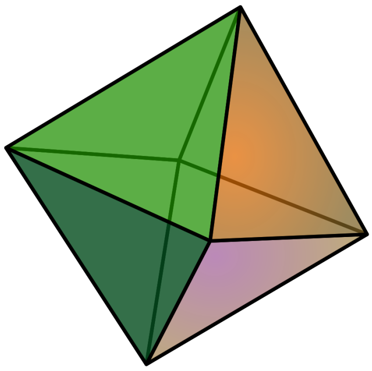

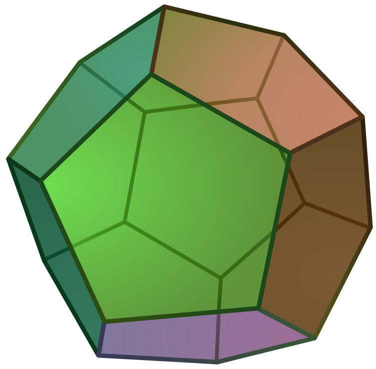

Figure 9.2: *Two 3-dimensional polytopes: The octahedron (left) and the dodecahedron (right)*

Vertex-edges graphs of 3-dimensional polytopes are well-understood, due to the following classical result. Ziegler's book has a full lecture about it [4, Lecture 4].

Theorem 9.1 (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 such graphs before and have shown that they have a unique (combinatorial) embedding in the plane; see Theorem 2.26; here, we see that they also have a geometric embedding as a polytope in $\mathbb{R}^3$. In particular, a polytope in $\mathbb{R}^3$ with n vertices has at most $3n - 6$ edges and $2n - 4$ faces, by Corollary 2.5.

Figure 9.3 shows the vertex-edge graph of the octahedron, drawn as a plane straight line graph (PSLG).

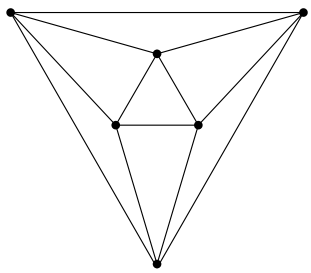

Figure 9.3: *The vertex-edge graph of the octahedron*

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

[1] By User:Stannered - Vectorisation of Image:Octahedron.jpg, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=1742116

[2] By DTR, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=2231561

polytope with n vertices have? Is it still $O(n)$, as for $d = 2, 3$? To discuss this, we first have to define polytope "edges" formally—our intuition unfortunately stops in $\mathbb{R}^3$. In fact, the vertices and edges are the 0-dimensional and 1-dimensional *faces* of the polytope.

9.1 Faces of Polytopes

Let $P \subset \mathbb{R}^d$ and $\mathcal{P} = \mathrm{conv}(P)$ a polytope. The dimension $\dim(\mathcal{P})$ of $\mathcal{P}$ is the dimension of the affine hull $\mathrm{aff}(P)$. If $\dim(\mathcal{P}) = d$, $\mathcal{P}$ is called *full-dimensional*.

Definition 9.2. *Let $\mathcal{P} \subset \mathbb{R}^d$ be a polytope. $F \subset \mathbb{R}^d$ is a* face *of $\mathcal{P}$ if there exists a hyperplane*

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

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

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

Here, h^+ is the closed positive halfspace *associated with the hyperplane h which is said to* support *the face F.*

Note that after multiplying all h_i with -1, the negative halfspace is the positive halfspace of some other hyperplane, so considering only positive halfspaces in the definition is no loss of generality. The way you should think about it is that a face is the intersection of $\mathcal{P}$ with some hyperplane "tangent" to $\mathcal{P}$. Figure 9.4 illustrates this notion.

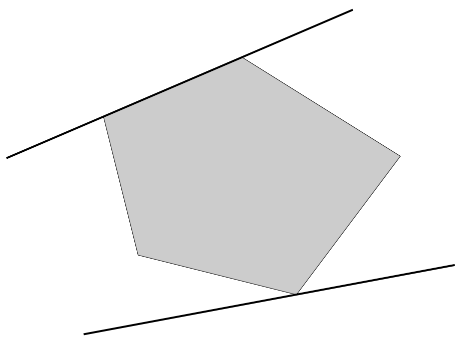

Figure 9.4: *Two faces (an edge and a vertex) with supporting hyperplanes.*

The *dimension* of a face is the dimension of its affine hull. A face of dimension k is called a k-face. There are established names for faces of certain dimensions.

- 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*.

By $V(\mathcal{P})$, we denote the set of vertices of a polytope $\mathcal{P}$.

If $\mathcal{P}$ is full-dimensional, the affine hull of a facet has dimension $d-1$ (is a hyperplane), from which we conclude the following.

Remark 9.3. *Every facet of a full-dimensional polytope has a unique supporting hyperplane.*

Exercise 9.4. *Show that any ridge is incident to exactly 2 facets.*

There are two degenerate hyperplanes obtained from setting $h_1 = \cdots = h_d = 0$.[3] If $h_{d+1} = 0$, we have $h = h^+ = \mathbb{R}^d$, and if $h_{d+1} < 0$ we get $h = \emptyset, h^+ = \mathbb{R}^d$. It follows that $\mathcal{P}$ is a face of itself, and that $\emptyset$ is also a face (by convention, its dimension is -1).

As an illustration, consider the octahedron; see Figure 9.2 (left). It has 8 facets, 12 edges (which are also ridges), and 6 vertices. The dodecahedron in Figure 9.2 (right) has 12 facets, 30 edges/ridges, and 20 vertices.

The careful reader might have noticed that previously (Definition 5.7), we have defined a vertex of $\mathcal{P} = \mathrm{conv}(P)$ differently, namely as an extreme point of P, one that is not a convex combination of the others. We should make sure that Definition 9.2 agrees with this, up to the formal subtlety that according to the latter, a vertex is not a point but a singleton set (we will later sweep this under the rug, but it is good to have talked about it once). The proof also shows how to argue with supporting hyperplanes.

Lemma 9.5. *Let $\mathcal{P} = \mathrm{conv}(P)$ be a polytope. Then p is an extreme point of P if and only if $\{p\}$ is a 0-face of $\mathcal{P}$.*

Proof. If p is an extreme point, then the disjoint sets $C = \{p\}$ and $D = \mathrm{conv}(P \setminus \{p\})$ are both closed and bounded convex sets, hence compact. By the Separation Theorem 5.16, there is a hyperplane that separates $\{p\}$ from $\mathrm{conv}(P \setminus \{p\})$. In formulas, there exist hyperplane parameters $h_1, \ldots, h_{d+1}$ such that

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

[3] In defining a hyperplane in Section 1.2, we haven't allowed for such degenerate hyperplanes, but here we need them.

After decreasing h_{d+1} and restricting to $q \in P \setminus \{p\} \subset \text{conv}(P \setminus \{p\})$, we get

$$\sum_{i=1}^{d} h_i p_i = h_{d+1}, \tag{9.6}$$

$$\sum_{i=1}^{d} h_i q_i > h_{d+1}, \quad q \in P \setminus \{p\}. \tag{9.7}$$

We claim that the hyperplane $h = \{x \in \mathbb{R}^d : \sum_{i=1}^{d} h_i x_i = h_{d+1}\}$ supports $\{p\}$, so $\{p\}$ is a 0-face. Indeed, p is on the hyperplane by (9.6), and any other point $p' \in \mathcal{P}$ is a convex combination $p' = \sum_{q \in P} \lambda_q q$, $\sum_{q \in P} \lambda_q = 1$, with $\lambda_q \geqslant 0$ for all q, and $\lambda_r > 0$ for some $r \neq p$. Then (9.6) and (9.7) yield

$$\sum_{i=1}^{d} h_i p'_i = \sum_{i=1}^{d} h_i \sum_{q \in P} \lambda_q q_i = \sum_{q \in P} \lambda_q \sum_{i=1}^{d} h_i q_i > \sum_{q \in P} \lambda_q h_{d+1} = h_{d+1}.$$

For the inequality, we used $\lambda_r > 0$ and $r \neq p$. Hence, $h \cap \mathcal{P} = \{p\}$ and $\mathcal{P} \subset h^+$.

For the other direction, suppose that p is not an extreme point, meaning that p is a convex combination of some other points $Q \subset P \setminus \{p\}$, $p = \sum_{q \in Q} \lambda_q q$, $\sum_{q \in Q} \lambda_q = 1$, with $\lambda_q > 0$ for all $q \in Q$. Any candidate for a hyperplane supporting $\{p\}$ must satisfy $h \cap \mathcal{P} \supseteq \{p\}$ and $Q \subset h^+$. If $h_1, \ldots, h_{d+1}$ are the parameters of h, this translates to

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

But this means that the inequality in the previous equation is in fact an equality, and as a consequence, h does not only contain p but also all points in Q. Hence, there is no hyperplane supporting $\{p\}$, so $\{p\}$ is not a 0-face. □

In a similar fashion, we could also convince ourselves that the concept of a 1-face indeed agrees with what we previously called an edge (a line segment connecting two vertices); but we refrain from doing so, hoping that the previous proof has provided enough credibility in this direction.

9.1.1 Faces and vertex sets

Here is an important fact: Every face F of a polytope $\mathcal{P}$ is a polytope itself whose vertices are exactly the vertices of $\mathcal{P}$ contained in F [4, Proposition 2.3. (i)]. Hence, a face is the convex hull of the vertices that it contains and is in particular uniquely identified by these vertices. As a consequence, if $\mathcal{P}$ has n vertices, then $\mathcal{P}$ has at most 2^n faces.

In the octahedron, each facet is a triangle, and it is a polytope (the convex hull of its three vertices). Each edge is the convex hull of its two endpoints and as such a polytope itself.

Exercise 9.8. *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 $k = 0, \ldots, \dim(\mathcal{P}) - 1$.*

This means, the total number of faces (excluding $\emptyset$ and $\mathcal{P}$) is at most

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

which is asymptotically (for $n \to \infty$) substantially less than 2^n. Also, 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. But we know that for $d = 2, 3$, this is a gross overestimate for large n, because in these dimensions, there are only $O(n)$ many edges.

Here is another very important property of faces [4, Proposition 2.3 (ii)].

Lemma 9.9. *Let F, G be two faces of a polytope $\mathcal{P}$. Then $F \cap G$ is also a face of $\mathcal{P}$.*

It also follows from the previous discussion that $V(F \cap G) = V(F) \cap V(G)$.

9.1.2 The Euler-Poincaré formula

For polytopes in $\mathbb{R}^3$, Eulers 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 9.10 (Euler-Poincaré formula). *For every d-dimensional polytope we have*

$$f_0 - f_1 + \ldots + (-1)^{d-1} f_{d-1} = 1 - (-1)^d.$$

For a proof of this formula, we refer to [4], Corollary 8.17.

Exercise 9.11. *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 $f_3 \geqslant f_0$.*

9.2 Polyhedra and the Main Theorem

We already know that a polytope can be written as an intersection of halfspaces, see Definition 5.7 and Theorem 5.19. Of particular interest for us is the finite case.

Definition 9.12. *A* polyhedron *is the intersection of finitely many halfspaces in $\mathbb{R}^d$.*

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); see Figure 9.5 for an example in $\mathbb{R}^2$.

The faces can be defined for a polyhedron $\mathcal{P}$ 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 9.5 has 4 vertices and 5 edges (= facets), two of them unbounded.

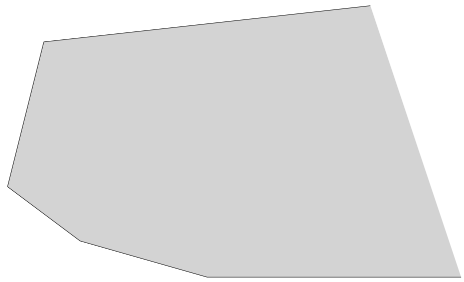

Figure 9.5: *An (unbounded) polyhedron in $\mathbb{R}^2$ (intersection of 5 halfspaces)*

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. This is indeed true in any dimension! Vice versa, every polytope is in particular a polyhedron. So polytopes and bounded polyhedra are the same objects. This is arguably the most fundamental result in polytope theory, and for this reason, Ziegler calls it the *Main Theorem* [4, Theorem 1.1].

Theorem 9.13 (Main Theorem). *A subset $\mathcal{P} \subset \mathbb{R}^d$ is the convex hull of a finite set of points $P \subset \mathbb{R}^d$ (a $\mathcal{V}$-polytope) if and only if $\mathcal{P}$ is a bounded intersection of finitely many halfspaces (an $\mathcal{H}$-polytope).*

Exercise 9.14. *Let $\mathcal{P} = \cap_{i=1}^m h_i^+$ be a full-dimensional polytope, represented as the intersection of m halfspaces $h_1^+, \ldots, h_m^+$, according to the Main Theorem. 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 [4, 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 $\mathcal{P} = \cap_{i=1}^m h_i^+$, where each h_i supports one of the facets. Hence, in the same way that non-extreme points are redundant in defining a $\mathcal{V}$-polytope, hyperplanes not supporting facets are redundant in defining an $\mathcal{H}$-polytope.

Corollary 9.15. *Let $\mathcal{P}$ be a full-dimensional polytope. Then every point p in the boundary of $\mathcal{P}$ is contained in some facet.*

Proof. We have $\mathcal{P} = \cap_{h \in H} h^+$ for the facet-supporting hyperplanes $h \in H$. Then p must be contained in at least one of these hyperplanes h, otherwise, a small neighborhood around p would still be in $\mathcal{P}$. So, $p \in \mathcal{P} \cap h =: F$, the facet supported by h. □

9.3 Examples

Let's look at two families of higher-dimensional polytopes that are the natural generalizations of two other platonic solids, the *cube*[4] and the *tetrahedron*[5]; see Figure 9.6.

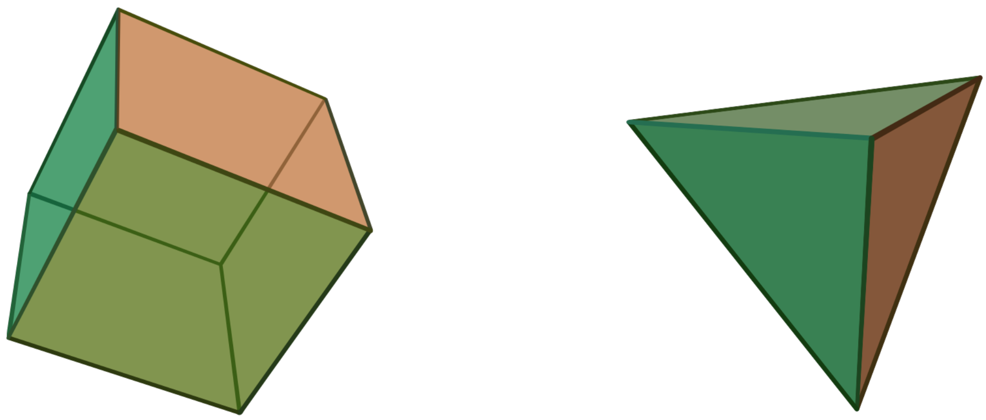

Figure 9.6: *The cube (left) and the tetrahedron (right)*

9.3.1 Hypercubes

Given some dimension $d \geqslant 1$, the d-dimensional *hypercube* is the set

$$C_d = \{x \in \mathbb{R}^d : -1 \leqslant x_i \leqslant 1, i = 1, \dots, d\}.$$

Formally, C_d is a polyhedron, described as the intersection of $2d$ halfspaces. But as C_d is bounded, the Main Theorem guarantees that C_d is a polytope. It has at most $2d$ facets by Exercise 9.14, but it is easy to see that it has exactly $2d$ facets (try to make the argument!). The next exercise is about the vertices of C_d.

Exercise 9.16. *Prove that C_d has 2^d vertices. What are they?*

9.3.2 Simplices

A k-*simplex* is the convex hull of $k + 1$ *affinely independent* points in $\mathbb{R}^d$. Note that we must have $k \leqslant d$, since more than $d + 1$ points cannot be affinely independent.

Exercise 9.17. *Prove that a k-simplex has the largest possible number of faces, namely 2^{k+1}, meaning that for every subset of vertices, there is a face with exactly these vertices.*

[4] By User:DTR - Vectorisation of Image:Hexahedron.jpg, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=2231470

[5] By !Original:Kjell AndréVector: DTR - Vectorisation of Tetrahedron.jpg, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=2231463

9.4 The graph of a polytope

For any polytope $\mathcal{P}$ in $\mathbb{R}^d$, its vertices and edges form an undirected graph $G(\mathcal{P})$, sometimes also called the *1-skeleton* of $\mathcal{P}$. For 3-dimensional polytopes , these graphs are well understood: recall that Steinitz' theorem states that a graph G is the graph of a 3-dimensional polytope if and only if it is planar and 3-connected. In higher dimensions, the graph of a d-dimensional polytope has minimum degree at least d.

These graphs are very interesting from a computational viewpoint, as we'll briefly explain here, without going into details. Consider an instance of *linear programming*: we want to maximize $c^T x$ subject to $Ax \leqslant b$. The inequalities $Ax \leqslant b$ define a finite set of halfspaces, whose intersection is the polyhedron of feasible regions. Let us assume that this region is non-empty and bounded. Then, by the Main Theorem, the feasible solutions form a polytope $\mathcal{P}$. Let g_0 be the maximum value that $c^T x$ can attain on $\mathcal{P}$. Then $c^T x = g_0$ defines a hyperplane whose intersection with $\mathcal{P}$ is the set of optimal solutions. In particular, the set of optimal solutions is a face of $\mathcal{P}$. Further, $c^T x$ defines directions on the edges of $G(\mathcal{P})$ where we orient v towards w whenever $c^T w > c^T w$. Clearly this graph is acyclic. Further, every sink is an optimal solution (Exercise 9.18). Thus, one way to find an optimal solution is to transverse the graph $G(\mathcal{P})$ along its 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*.

Exercise 9.18. *Let $\mathcal{P}$ be a polytope. Consider the directed graph induced by the linear functional $c^T x$. Show that for every vertex $v \in \mathcal{P}$ that is not an optimal solution, 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. This was conjectured by Warren M. Hirsch in 1957, whose conjecture is now known as the *Hirsch conjecture*: for a d-dimensional polytope $\mathcal{P}$ with n facets, the diameter of the graph $G(\mathcal{P})$ is at most $n - d$.

This conjecture was disproven in 2010 by Francisco Santos, who constructed a 43-dimensional polytope with 86 facets whose graph has diameter larger than 43 [2]. However, the *polynomial Hirsch conjecture*, which states that for a d-dimensional polytope $\mathcal{P}$ with n facets, the diameter of the graph $G(\mathcal{P})$ is at most polynomial in n, is still open.

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

Theorem 9.19 (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 $|V| \geqslant d+1$. We want to show that deleting a subset $S \subset V$ of $d-1$ vertices does not disconnect $G(\mathcal{P})$. We distinguish two cases.

Case 1. Assume that all vertices in S are contained in a proper face F. Let $h = \{x \in \mathbb{R}^d : \sum_{i=1}^{d} h_i x_i = h_{d+1}\}$ be a hyperplane that supports F. In particular, h_{d+1} is the smallest value that $\sum_{i=1}^{d} h_i x_i$ can attain on $\mathcal{P}$. Denote by $g_{d+1} > h_{d+1}$ the largest value that $\sum_{i=1}^{d} h_i x_i$ can attain on $\mathcal{P}$ and let F_0 be the face supported by the hyperplane $g = \{x \in \mathbb{R}^d : \sum_{i=1}^{d} h_i x_i = g_{d+1}\}$. For every vertex $v \in V$, let $f(v)$ denote the unique value such that v lies on the hyperplane $\{x \in \mathbb{R}^d : \sum_{i=1}^{d} h_i x_i = f(v)\}$. Now, by Exercise 9.18, every vertex $v \in V$ is either in F_0 or it has a neighbor $w \in V$ with $f(w) > f(v)$. In particular, w lies in $V \setminus S$. Thus, from every vertex v there is a path to a vertex $u \in F_0$ along which f strictly increases and which avoids S. Finally, The graph on F_0 is connected by induction on d.

Case 2. Assume that the vertices in S are not contained in a proper face. Let $h = \{x \in \mathbb{R}^d : \sum_{i=1}^{d} h_i x_i = h_{d+1}\}$ be a hyperplane that contains S and at least one more vertex $v_0 \in V \setminus S$. Such a hyperplane exists, as every set of d points is contained in a hyperplane. Let $g_{\max}$ and $g_{\min}$ denote the maximum and minimum value that $\sum_{i=1}^{d} h_i x_i$ can attain on $\mathcal{P}$, respectively, and let $F_{\max}$ and $F_{\min}$ be the corresponding faces. By induction on d, the graphs $G(F_{\max})$ and $G(F_{\min})$ are connected. Similar to Case 1, each vertex $v \in V \setminus S$ is either connected by an increasing path to $F_{\max}$ or by a decreasing path to $F_{\min}$, and these paths avoid S. Finally, v_0 is connected to both $F_{\max}$ and $F_{\min}$. □

9.5 Polytope Structure

In this section, we will present some more advanced properties of polytopes, mostly without proofs, as these would take us too far away from our actual subject (geometry). But all of the following is classical material and can be found in full detail for example in Ziegler's book [4].

9.5.1 The face lattice

The face lattice of a polytope $\mathcal{P}$ is the partial order on the set $\mathcal{F}$ of faces of $\mathcal{P}$, ordered by inclusion. In this partial order, we have $F \leqslant G$ if $F \subset G$. We say that $F < G$ if $F \leqslant G$ and $F \neq G$. Partial orders are usually drawn as *Hasse diagrams* where larger elements are higher up, smaller elements are further down, and *cover* relations ($F < G$ but there is no H such that $F < H < G$) are drawn as connections between the elements. For example, the face lattice of the 3-dimensional cube is depicted in Figure 9.7.

What makes this partial order a *lattice* [4, Theorem 2.7] is the fact that for any two faces F and G, there is (a) a unique inclusion-maximal face E such that $E \subset F, G$ (the *meet* of F and G), and (b) a unique inclusion-minimal face H such that $F, G \subset H$ (the *join* of F and G). The meet of F and G is $F \cap G$ which is also a face by Lemma 9.9, with vertex set $V(F) \cap V(G)$. It may be tempting to believe that the join of F and G is the face with vertex set $V(F) \cup V(G)$, but this is not true in general, as there may be no such face. We already see this in the cube. The join of two incident edges (for example 12 and 13) is a face with four vertices (1234).

Exercise 9.20. *While the existence of a meet of* F *and* G *easily follows from the fact*

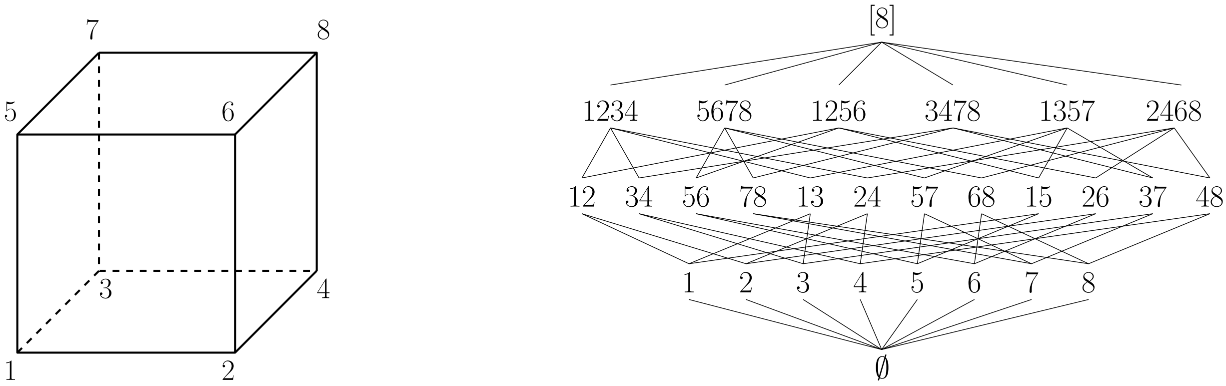

Figure 9.7: *The cube (left) and its face lattice (right). Faces are named with the labels of their vertices.*

that $F \cap G$ is a face (think about the precise argument!), the existence of a join is less clear. However, this already follows from the existence of meets. To make this formal (and at the same time review the poset terminology), you are asked to prove the statement for general posets.

A poset *(partially ordered set) is a pair $(\mathcal{F}, \leqslant)$ where $\leqslant$ is a relation over $\mathcal{F}$ that is* reflexive *($F \leqslant F$ always holds),* antisymmetric *($F \leqslant G$ and $G \leqslant F$ implies $F = G$) and* transitive *($F \leqslant G$ and $G \leqslant H$ implies $F \leqslant H$). By $F < G$ we mean $F \leqslant G$ and $F \neq G$. A* maximal element *in $(\mathcal{F}, \leqslant)$ is an element 1 such that there is no element $1 < H$. Similarly, a* minimal element *is an element 0 such that there is no element $E < 0$.*

A largest lower bound *of F and G is an element E such that $E \leqslant F$ and $E \leqslant G$, but no element $E' > E$ has this property. If F and G have a unique largest lower bound, we call it the* meet *of F and G. Similarly, a* smallest upper bound *of F and G is an element H such that $F \leqslant H$ and $G \leqslant H$, but no element $H' < H$ has this property. If F and G have a unique smallest upper bound, we call it the* join *of F and G.*

Now for the actual exercise: Let $(\mathcal{F}, \leqslant)$ 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!

The face lattice represents the combinatorial information contained in a polytope. Two polytopes are called *combinatorially equivalent* if they have isomorphic face lattices [4, Section 2.2]. Combinatorially equivalent polytopes may geometrically look quite different. For examples, any two triangles in the plane are combinatorially equivalent, but one of them could be a nice (Delaunay) triangle, and the other one a long and skinny triangle.

9.5.2 Polarity

For every polytope $\mathcal{P} \ni 0$, there is a polytope $\mathcal{P}^\triangle \ni 0$ (the polar polytope) whose face lattice is that of $\mathcal{P}$, turned upside down [4, Theorem 2.11]. This means, vertices of $\mathcal{P}$ correspond to facets of $\mathcal{P}^\triangle$, edges of $\mathcal{P}$ to ridges of $\mathcal{P}^\triangle$, and so on.

If $\mathcal{P} = \mathrm{conv}(P)$, then $\mathcal{P}^\triangle = \cap_{p \in P} h_p^+$, where

$$h_p^+ = \{x \in \mathbb{R}^d : \sum_{i=1}^{d} p_i x_i \leqslant 1\}.$$

This is not a positive halfspace as in Definition 9.2, but after multiplying all numbers with -1, we arrive at an equivalent positive halfspace. It can be shown that $\mathcal{P}^{\triangle\triangle} = \mathcal{P}$.

Geometrically, going to the polar polytope corresponds to replacing a point (contributing to $\mathcal{P}$ as a convex hull) with a halfspace (contributing to $\mathcal{P}^\triangle$ as an intersection of halfspaces); see Figure 9.8.

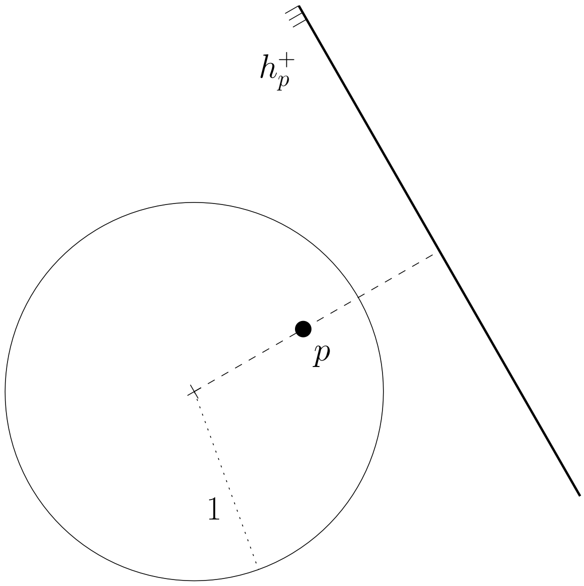

Figure 9.8: *The polar halfspace* h_p^+ *corresponding to a point* p *has distance* $1/\|p\|$ *from* 0 *and is perpendicular to the vector* p. *This operation is called* inversion at the unit sphere.

We can also "polarize" $\mathcal{P}$ if $0 \notin P$, by simply choosing the center of inversion as some point in the interior of $\mathcal{P}$. Depending on which point we choose, $\mathcal{P}^\triangle$ will look different but its combinatorial structure (face lattice) will always be the same.

We already know some pairs of mutually polar polytopes. Indeed, each platonic solid is polar to another one; see Figure 9.9.

For example, the dodecahedron has 12 facets (hence its name), 30 edges and 20

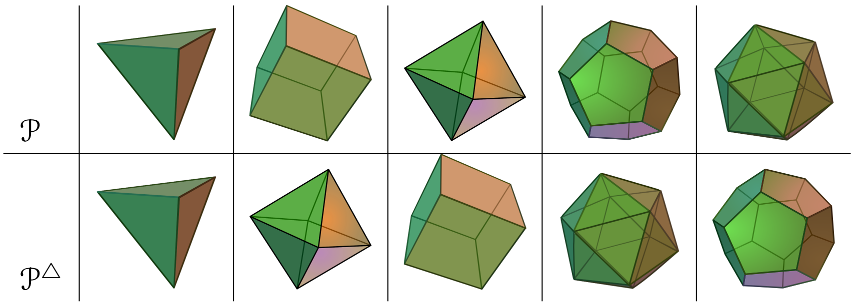

Figure 9.9: *Polarities among the platonic solids: the tetrahedron is polar to itelf (first column); cube and octahadreon are polar to each other (second and third column); dodehacedron and icosahedron are polar to each other (fourth and fifth column).*

vertices; it's polar, the icosahedron[6] has 20 facets (hence its name), 30 edges and 12 vertices.

Three of the platonic solids in Figure 9.9 generalize to polytopes in arbitrary dimension d, and we have already encountered two of these generalizations in Section 9.3: simplices and hypercubes. Simplices are polar to simplices, and hypercubes are polar to *cross-polytopes* that generalize the octahedron. The standard d-dimensional cross-polytope is the convex hull of the d unit vectors in $\mathbb{R}^d$ and their negatives, so it has 2d vertices (and 2^d facets).

9.6 Simplicial and Simple Polytopes

An important question about polytopes that we want to answer is the following:

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

We already know that for $d = 2$, the answer is n. For $d = 3$, it's at most $2n - 4$. In both cases, the bound is linear in the number of vertices. For general d, we get a bound of $O(n^d)$ from Exercise 9.8, but we see that already for $d = 2, 3$, this bound is an overestimate.

To address the above question (that we will fully answer only in a later chapter), it turns out that we can restrict our attention to *simplicial polytopes*. These are polytopes where all facets are $(d-1)$*-simplices*. For example, the octahedron in Figure 9.2 (left) is simplicial, since all its facets are triangles (2-simplices). The dodacehdron in Figure 9.2 (right) is not simplicial, since its facets are pentagons.

For a given 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"

[6] By User:DTR - Vectorisation of Image:Icosahedron.jpg by en:User:Cyp, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=2231553

by slightly perturbing its vertices (doing this randomly works with probability 1). What happens under such a perturbation is that each non-simplicial facet "breaks apart" and gets replaced by some simplicial facets (facets that are $(d-1)$-simplices). Facets that have been simplicial before remain simplicial.

Let's look at this for the cube $[0,1]^3$. Suppose that we move the two vertices $(0,0,0)$ and $(1,1,1)$ "slightly inwards" so that they become $(\varepsilon,\varepsilon,\varepsilon)$ and $(1-\varepsilon,1-\varepsilon,1-\varepsilon)$, respectively, for some small $\varepsilon > 0$, then we obtain the polytope in Figure 9.10.

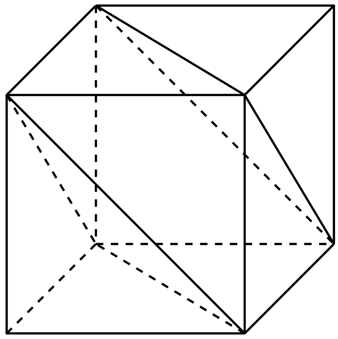

Figure 9.10: *Perturbing the cube vertices: by moving the lower left and the upper right vertex slightly inwards, each square facet breaks up into two triangles, and the resulting polytope is simplicial.*

Exercise 9.21. *What happens if we move the two vertices $(0,0,0)$ and $(1,1,1)$ "slightly outwards" so that they become $(-\varepsilon,-\varepsilon,-\varepsilon)$ and $(1+\varepsilon,1+\varepsilon,1+\varepsilon)$, respectively? Draw the resulting simplicial polytope!*

Similarly, for the dodecahedron, each pentagon facet gets replaced by three triangles when we slightly perturb the vertices. In particular, the number of facets increases under this perturbation, so any polytope with the largest number of facets for a given number of vertices is simplicial (seee [4, Lemma 8.24] for a formal statement and reference to a proof).

Simplicial polytopes are particularly nice. We have the following equivalent characterization [4, Proposition 2.16].

Theorem 9.22. *A polytope is simplicial if and only if every k-face has $k+1$ vertices, $0 \leqslant k \leqslant d-1$.*

In view of the fact that every face is the convex hull of its vertices, this just means that in a simplicial polytope, not only the facets, but *all* faces are simplices of the appropriate dimension.

Via polarity, our initial question can be reformulated as follows:

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

Here, the polytopes maximizing the count are the *simple* ones where every vertex is incident to d edges. As the polarity transform turns the face lattice upside down, a polytope is simple if and only if its polar polytope is simplicial. Checking Figure 9.9 again, we see that the tetrahedron is both simple and simplicial, the octahedron as well as the icosahedron are simplicial, and their duals— the cube and the dodecahedron—are simple.

Exercise 9.23. *Characterize all polytopes in* $\mathbb{R}^3$ *that are both simple and simplicial.*

9.7 Higher-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 flips are possible anymore. In this section, we want to discuss more systematically what the "lifted surface" actually is when we are done with the Lawson flips, meaning that the triangulation has become Delaunay. In fact, we want to do this for any dimension d.

We will give the big picture upfront, borrowing the very nice Figure 9.11 below from Hang Si [3].

Let us assume general position (no three points on a line, no four points on a circle). In this situation, the Delaunay triangulation of a planar point set is unique, see Corollary 6.19. An alternative way to look at it is the following: lift 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, when projected back to $\mathbb{R}^2$, give us the Delaunay triangulation; see Figure 9.11 (lower right part).

So the "lifted surface" after the Lawson flip algorithm has terminated is actually the *lower convex hull* of the lifted points. This also means that we can reduce the computation of the Delaunay triangulation to the computation of a convex hull in $\mathbb{R}^3$. This is what we will formally prove in this section, for general dimension d.

Figure 9.11 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 what is called the *farthest-point Delaunay triangulation*. It is in general not a triangulation of the point set, but only of the points on the convex hull. Each triangle in this triangulation is an "anti-Delaunay" triangle in the sense that its circumcircle contains *all* other points; see Exercise 9.40 below.

The left part of Figure 9.11 shows what happens if we lift the points not onto the unit paraboloid in $\mathbb{R}^3$ but in some arbitrary way. The convex hull of the lifted point set is still a polytope, and if the lifting is such that this polytope is simplicial, we can recover two triangulations in the plane, coming from projecting the lower facets back to $\mathbb{R}^2$, and from projecting the upper facets back to $\mathbb{R}^2$. Such triangulations are called *regular*; the (farthest-point) Delaunay triangulation is a specific regular triangulation.

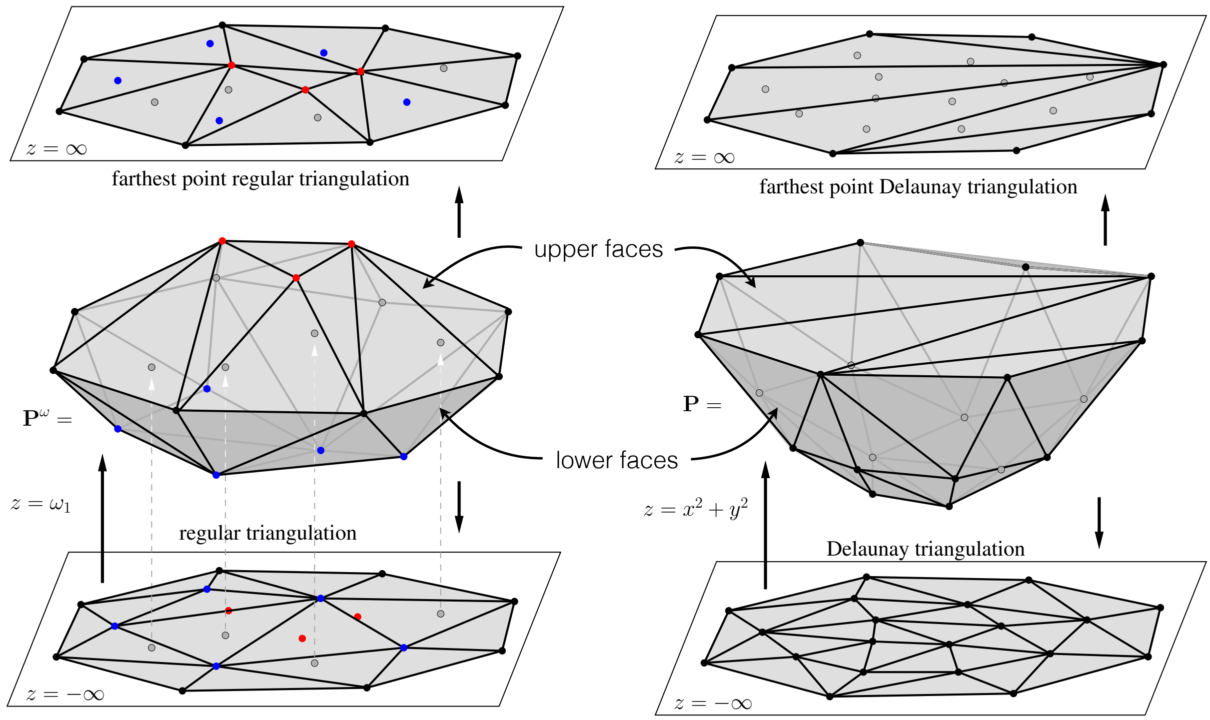

Figure 9.11: *Triangulations in* $\mathbb{R}^d$ *as projections of polytopes in* $\mathbb{R}^{d+1}$

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

Definition 9.24. *A* **triangulation** *of a finite point set* $P \subset \mathbb{R}^d$ *is a collection* $\mathcal{T}$ *of* d*-simplices, such that*

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

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

(3) for every distinct pair $T, U \in \mathcal{T}$*, the intersection* $T \cap U$ *is a face of both.*

Note that this also allows for $T \cap U = \emptyset$, since $\emptyset$ is a face of every polytope.

At this point, it is not even clear whether every point set in $\mathbb{R}^d$ has a triangulation if $d \geqslant 3$. For $d = 2$, we recover Definition 6.1, and also for $d = 1$, the general definition makes sense. A point set $\{a_1, a_2, \ldots, a_n\}$ in $\mathbb{R}^1$ (where we assume that $a_1 < a_2 < \cdots a_n$) has a unique triangulation $\mathcal{T} = \{[a_i, a_{i+1}] : 1 \leqslant i < n\}$.

We can also define Delaunay triangulations in the same way as before.

Definition 9.25. *A* **Delaunay triangulation** *of a finite point set* $P \subset \mathbb{R}^d$ *is a triangulation* $\mathcal{T}$ *of* P*, with the property that the circumsphere of every* d*-simplex in* $\mathcal{T}$ *is empty of points from* P*.*

What is the circumsphere of a d-simplex T? This is the unique sphere that contains all vertices $V(T)$. Before you can even ask whether such a sphere always exists, and why it is unique, *if* it exists, let us prove it.

Lemma 9.26. *Let* $S \subset \mathbb{R}^d$ *be a set of* $d+1$ *affinely independent points. Then there exists a unique sphere containing* S*.*

Proof. A sphere has a center $c \in \mathbb{R}^d$ and a real radius $r \geqslant 0$, and is formally defined as the set

$$\{x \in \mathbb{R}^d : \|x - c\| = r\}.$$

Squaring the condition, we are therefore looking for a (unique) point $c \in \mathbb{R}^d$ and a (unique) number r^2 such that

$$\|p - c\|^2 = r^2, \quad p \in S. \tag{9.27}$$

For a point $x \in \mathbb{R}^d$ (a column vector), let us denote by $x^\top$ its transpose (a row 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 can equivalently be written as

$$p^\top p = 2p^\top c + \underbrace{r^2 - c^\top c}_{=:\alpha}, \quad p \in S. \tag{9.28}$$

In still other words,

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

This system of equations is of the form $b = (2c^\top, \alpha)B$, where the entries of the row vector $b \in \mathbb{R}^{d+1}$ are the $p^\top p, p \in S$, and the columns of the $(d+1) \times (d+1)$ matrix B are the $\begin{pmatrix} p \\ 1 \end{pmatrix}, p \in S$. As the $p \in S$ are affinely independent, the columns of B are lineary independent (Proposition 5.3), and hence B is invertible. So there are is a unique $c \in \mathbb{R}^d$ and a unique $\alpha \in \mathbb{R}$ such that (9.28) holds, meaning that (9.27) holds with a unique c and unique $r^2 := \alpha + c^\top c$ which must then be nonnegative. □

Next we want to show that we can always find a unique Delaunay triangulation, assuming sufficiently general position. To prepare this, we first define the concept of a Delaunay simplex.

Definition 9.29. *Let* $P \subset \mathbb{R}^d$ *be a set of points in general position, meaning that no* $d+1$ *points lie on a common hyperplane, and no* $d+2$ *points lie on a common sphere. A simplex* $\mathrm{conv}(S)$ *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*.*

Here is the crucial insight. Generalizing the lifting map (Lemma 6.12), we can show that Delaunay simplices correspond to (lower) facets of a polytope in one dimension higher, namely the convex hull of the lifted points. For $p = (p_1, \ldots, p_d) \in \mathbb{R}^d$, we define the *lifted point*

$$\ell(p) = (p_1, \ldots, p_d, p^\top p) \in \mathbb{R}^{d+1}. \tag{9.30}$$

For $d = 2$, we recover the standard lifting map that lifts points in the plane to the unit paraboloid in $\mathbb{R}^3$, see Section 6.3.

Lemma 9.31. *Let $P \subset \mathbb{R}^d$ be in general position according to Definition 9.29, and let $\mathcal{P} = \mathrm{conv}(\ell(P))$ be the convex hull of the lifted points. Then $\mathcal{P}$ has vertex set $\ell(P) = \{\ell(p) : p \in P\}$.*

Proof. By definition, $V(\mathcal{P}) \subset \ell(P)$, so it remains to show that $\ell(p)$ is a vertex of $\mathcal{P}$ for all $p \in P$. To this end, consider the hyperplane

$$h = \{x \in \mathbb{R}^{d+1} : x_{d+1} - \sum_{i=1}^{d} 2p_i x_i = -p^\top p\}. \tag{9.32}$$

For any $q \in \mathbb{R}^d$, and with $x = \ell(q)$ we have

$$x_{d+1} - \sum_{i=1}^{d} 2p_i x_i + p^\top p = q^\top q - \sum_{i=1}^{d} 2p_i q_i + p^\top p = \|p - q\|^2 \begin{cases} = 0, & q = p, \\ > 0, & q \neq p \end{cases},$$

meaning that h contains precisely one element of $\mathcal{P}$, namely $\ell(p)$. Indeed, we have just shown all other candidate vertices $\ell(p'), p' \in P \setminus \{p\}$ to be strictly above, and this implies that *any* point in $\mathcal{P} \setminus \{\ell(p)\}$ is strictly above; see the arguments in the proof of Lemma 9.5. Hence, $h \cap \mathcal{P} = \{\ell(p)\}$ and $\mathcal{P} \subset h^+$, so h is supporting the face $\{\ell(p)\}$ of $\mathcal{P}$. □

Lemma 9.33. *Let $P \subset \mathbb{R}^d$ be in general position according to Definition 9.29, and let $\mathcal{P} = \mathrm{conv}(\ell(P))$ be the convex hull of the lifted points. Then $\mathcal{P}$ is a simplicial polytope in $\mathbb{R}^{d+1}$. Moreover, let $S \in \binom{P}{d+1}$. Then the following two statements are equivalent.*

(i) $\mathrm{conv}(S)$ *is a Delaunay simplex for* P.

(ii) $\mathrm{conv}(\ell(S))$ *is a lower facet of* $\mathcal{P}$*, where a lower facet is one that has a supporting hyperplane of the form* $h = \{x \in \mathbb{R}^{d+1} : \sum_{i=1}^{d+1} h_i x_i = h_{d+2}\}$ *with* $h_{d+1} = 1$.

Proof. We have already done most of the work in the proof of Lemma 9.26. Let $c \in \mathbb{R}^d$ and $r \in \mathbb{R}$ be center and radius of S's circumsphere, and let $\alpha = r^2 - c^\top c$. Also considering inequality variants of (9.27) and (9.28), we get that $\mathrm{conv}(S)$ is a Delaunay simplex for P if and only if

$$p^\top p = 2p^\top c + \alpha, \quad p \in S, \tag{9.34}$$
$$p^\top p > 2p^\top c + \alpha, \quad p \in P \setminus S. \tag{9.35}$$

An equality for $p \in P \setminus S$ cannot happen by general position. Defining $h_i = -2c_i$ for $i = 1, \ldots, d$, $h_{d+1} = 1$ and $h_{d+2} = \alpha$, this yields

$$\sum_{i=1}^{d+1} h_i \ell(p)_i = h_{d+2}, \quad p \in S, \tag{9.36}$$

$$\sum_{i=1}^{d+1} h_i \ell(p)_i > h_{d+2}, \quad p \in P \setminus S. \tag{9.37}$$

Equivalently, the hyperplane $h = \{x \in \mathbb{R}^{d+1} : \sum_{i=1}^{d+1} h_i x_i = h_{d+2}\}$ supports the face $\text{conv}(\ell(S))$ spanned by $d+1$ (affinely independent) vertices $\ell(S)$, so $\text{conv}(\ell(S))$ is a lower facet of $\mathcal{P}$.

We can also go backwards from (9.36) and (9.37) to (9.34) and (9.35) by defining $c_i = -h_i/2$ and $\alpha = h_{d+2}$, so if $\text{conv}(\ell(S))$ is a (lower) facet, then the points in S are on a common (empty) sphere. By general position, $|S| = d+1$, and $\text{conv}(S)$ is a Delaunay simplex. This in particular also shows that $\mathcal{P}$, the convex hull of the lifted points, is simplicial. □

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

Theorem 9.38. *Let $P \subset \mathbb{R}^d$ be in general position according to Definition 9.29, and let $\mathcal{P} = \text{conv}(\ell(P))$ be the convex hull of the lifted points (a polytope in $\mathbb{R}^{d+1}$). Then the set*

$$\mathcal{T} = \{\text{conv}(S) : S \in \binom{P}{d+1}, \text{conv}(\ell(S)) \textit{ is a lower facet of } \mathcal{P}\}$$

is the unique Delaunay triangulation of P.

Proof. If $\mathcal{T}$ is indeed a triangulation, then it is a Delaunay triangulation by Lemma 9.33 and in fact contains *all* possible Delaunay simplices; so another Delaunay triangulation can only contain less simplices, but then it has "holes" and is not even a triangulation.

It remains to prove that $\mathcal{T}$ is a triangulation, so let's look at the three required properties in Definition 9.24. For property (1)—the simplices cover exactly $\text{conv}(P)$—let $q \in \text{conv}(P)$. We need to construct a simplex $\text{conv}(S) \in \mathcal{T}$ containing q. Choose $t \in \mathbb{R}$ minimal such that the vertically lifted point $(q, t) \in \mathbb{R}^{d+1}$ is in $\mathcal{P}$. As $\mathcal{P}$ intersects the vertical line through q in a closed interval (as a consequence of $\mathcal{P}$ being closed and convex), t is attained at the lowest point of this interval. We have to argue that this interval is nonempty, but this holds, since $q \in \text{conv}(P)$ means that *some* (q, t) is in $\text{conv}(\ell(P)) = \mathcal{P}$.

The lifted point (q, t) is on the boundary of $\mathcal{P}$ and hence contained in one or more facets, by Corollary 9.15. One of these facets must be a lower facet $\text{conv}(\ell(S))$ (Exercise 9.39); as the lifted point is in $\text{conv}(\ell(S))$, the original point is in $\text{conv}(S) \in \mathcal{T}$. This shows that $\text{conv}(P) \subset \bigcup_{T \in \mathcal{T}} \text{conv}(T)$; to complete the proof of (1), it remains to prove the inclusion $\supset$ which follows from $\text{conv}(S) \subset \text{conv}(P)$ for all $S \in \binom{P}{d+1}$.

For property (2)—$P = \bigcup_{T \in \mathcal{T}} V(T)$—we show that every vertex $\ell(p), p \in P$ is the vertex of some lower facet of $\mathcal{P}$, implying that p is the vertex of some Delaunay simplex in $\mathcal{T}$. We first observe that if we choose t minimal such that $(p, t) \in \mathcal{P}$, we obtain $(p, t) = \ell(p)$; indeed, $t = p^\top p$ yields $\ell(p)$, but no smaller t leads to a point in $\mathcal{P}$, as such a (p, t) is outside the unit paraboloids "bowl" $\mathcal{U} = \{x \in \mathbb{R}^{d+1} : x_{d+1} \geqslant \sum_{i=1}^{d} x_i^2\}$ while $\mathcal{P} \subset \mathcal{U}$ by convexity of the bowl $\mathcal{U}$.

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.

Property (3)—the intersection of any two simplices T, U in $\mathcal{T}$ is a face of both—follows from polytope properties: let $T' = \ell(T)$ and $U' = \ell(U)$ be the (lower) facets of $\mathcal{P}$ corresponding to the Delaunay simplices T and U. The intersection $T' \cap U'$ is a face of $\mathcal{P}$ by Lemma 9.9, with vertex set $\ell(V(U)) \cap \ell(V(T)) = \ell(V(U) \cap V(T))$. Hence $T' \cap U' = \text{conv}(\ell(V(U) \cap V(T)))$ which implies that $T \cap U = \text{conv}(V(T) \cap V(U))$. This is a face of both simplices T and U, since every subset of vertices of a simplex defines a face (Exercise 9.17). □

Exercise 9.39. *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 9.40. *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 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 $P \setminus V(T)$:*

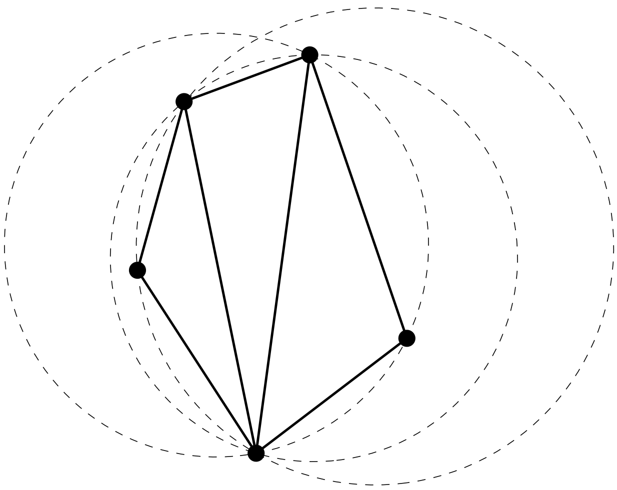

Prove that P has a unique farthest-point Delaunay triangulation; Figure 9.11 provides the intuition. The name comes from the fact that in the plane, the farthest-point Delaunay triangulation is dual to the farthest-point Voronoi diagram, *the subdivision of the plane into regions with the same* farthest *point.*

9.8 Complexity of 4-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 $O(n)$. We have also seen that for $d = 4$, the complexity is bounded by $O(n^4)$ (Exercise 9.8). But can we actually have superlinear complexity $\Omega(n)$ for $d = 4$, or does the "nice" behavior in dimensions $d = 2, 3$ continue?

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

Theorem 9.41. *For every even natural number $n \geqslant 4$, there exists a 4-dimensional simplicial polytope with n vertices and at least $(\frac{n}{2} - 1)^2 = \Theta(n^2)$ facets.*

Moreover, this polytope also has $\Theta(n^2)$ edges which is asymptotically maximal since Exercise 9.8 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 [4, Corollary 0.8]. This may be somewhat counter-intuitve, as it seems to require that many edges "go through" the polytope which they obviously cannot. On the other hand, 4 dimensions are counterintuitive *per se*, so let's not worry to 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 $(\frac{n}{2} - 1)^2$ Delaunay simplices. By Lemma 9.33, the convex hull of the lifted point set $\ell(P)$ is a 4-dimensional simplicial polytope with at least $(\frac{n}{2} - 1)^2$ (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 if four points $p, q, r, s \in P$ are chosen such that p, q are consecutive along ℓ_1, and r, s are consecutive along ℓ_2, then $\text{conv}(\{p, q, r, s\})$ is a Delaunay simplex. As there are $(n/2 - 1)^2$ ways to choose p, q, r, s in this way, P has the required number of Delaunay simplices.

It remains to prove the claim. For this, we refer to the (2-dimensional cartoon) Figure 9.12.

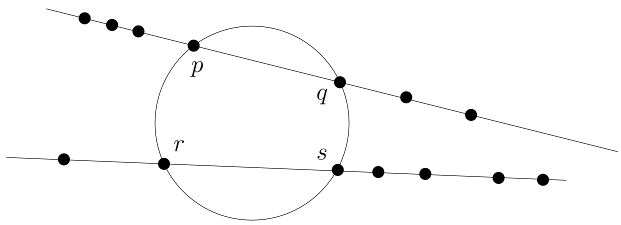

Figure 9.12: *Proof of Theorem 9.41*

As a consequence of the lines being skew, p, q, r, s are affinely independent and hence have a unique circumsphere. The line ℓ_1 intersects this sphere in exactly the points p and q, and as p, q are consecutive along ℓ_1, there is no point of P_1 inside the sphere. For the same reason, no point of P_2 is inside, so the sphere is empty, and $\text{conv}(\{p, q, r, s\})$ is a Delaunay simplex. □

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 9.41 is asymptotically best possible. We will not prove this here but in a later chapter give a (tight) upper bound on the number of facets that a d-dimensional polytope with n vertices can have.

9.9 Higher-dimensional Voronoi diagrams

We can also obtain the Voronoi diagram of a finite point set $P \subset \mathbb{R}^d$ from the facets of a *polyhedron* in dimension $d + 1$. In fact, Theorem 8.17 has already done this for $d = 2$, without explicitly mentioning polyhedra. Here, we simply reprove this theorem for general d; no new ideas appear, so the reader is invited to consider this section as a repetition of Section 8.4, but formulated in the language of polytopes and polyhedra, and replacing "2" by "d".

Let's start by generalizing Voronoi regions to higher-dimensions which is a straightforward adaptation of Definition 8.3.

Definition 9.42. *Let $P \subset \mathbb{R}^d, |P| = n$. For $p \in P$ denote the* **Voronoi cell** *$V_P(p)$ of p by*

$$V_P(p) := \left\{q \in \mathbb{R}^2 : \|q - p\| \leqslant \|q - p'\| \text{ for all } p' \in P\right\}.$$

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 9.43. *Let $P \subset \mathbb{R}^d, |P| = n$. For $p \in P$, let h_p be the hyperplane*

$$h_p = \{x \in \mathbb{R}^{d+1} : x_{d+1} - \sum_{i=1}^{d} 2p_i x_i = -p^\top p\},$$

and let $\mathcal{P} = \cap_{p \in P} h_p^+$ (a polyhedron in $\mathbb{R}^{d+1}$). Then all the $h_p, p \in P$, are (lower) facet-supporting hyperplanes of $\mathcal{P}$.

Moreover, let $q \in \mathbb{R}^d$, and choose $t \in \mathbb{R}$ minimal such that the vertically lifted point $(q, t) \in \mathbb{R}^{d+1}$ is in $\mathcal{P}$. Then the following two statements are equivalent.

(i) $q \in V_P(p)$.

(ii) $(q, t) \in h_p$, *meaning that (q, t) is in the facet of $\mathcal{P}$ supported by h_p.*

This says that the Voronoi cell $V_P(p)$ is obtained as the projection of the facet $\mathcal{P} \cap h_p$ of $\mathcal{P}$ to $\mathbb{R}^d$. Hence, if we project all the facets of $\mathcal{P}$ to $\mathbb{R}^d$, we obtain the Voronoi diagram of P, the subdivsion of $\mathbb{R}^d$ into regions where the closest point among all points in P is the same.

Figure 9.13, borrowed from the book by Joswig and Theobald [1, Page 87], visualizes this for $d = 3$.

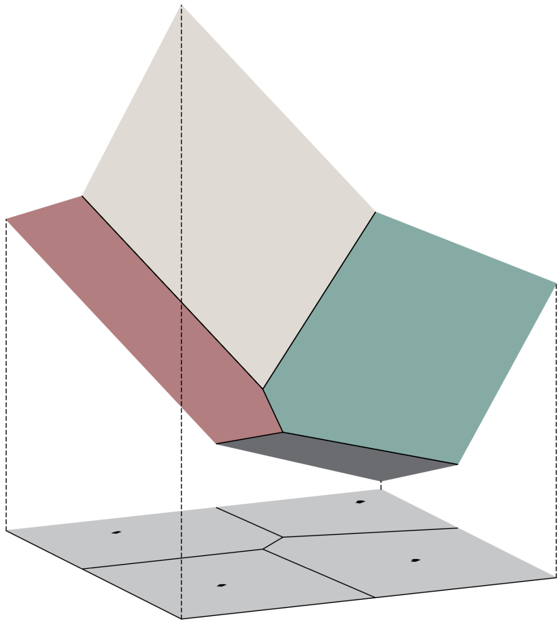

Figure 9.13: *A view from the "outside" on (a part of) the polyhedron $\mathcal{P} \subset \mathbb{R}^3$ in Theorem 9.43, and (a part of) the Voronoi diagram resulting from the projections of the facets to $\mathbb{R}^2$*

Proof of Theorem 9.43. 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 Section 9.2. Actually, we have employed the hyperplane h_p before in (9.32) to show that $\ell(p) \in h_p$ but $\ell(q) \in h_p^+ \setminus h_p$ for all $q \neq p$ where ℓ is the lifting map (9.30). Applying this argument for $p' \in P \setminus \{p\}$ and $q = p$, we see that $\ell(p) \in h_{p'}^+ \setminus h_{p'}$ for all $p' \in P \setminus \{p\}$, so $\ell(p)$ is in fact in the interior of $\cap_{p' \in P \setminus \{p\}} h_{p'}^+$ but on the boundary of $\cap_{p' \in P} h_{p'}^+$, so h_p is not redundant.

For the equivalence of (i) and (ii), we claim that the vertical distance of $\ell(q)$ to h_p is precisely $\|q - p\|^2$ (see Lemma 8.15 and Figure 8.7 for the 2-dimensional case). As $\ell(q)$ is above all the h_p (see first part of the proof), it follows that a hyperplane h_p is vertically closest to $\ell(q)$ and hence highest at q if and only if $q \in V_P(p)$. As (q, t) is in the highest hyperplane at q (that's where $\mathcal{P}$ "starts" when we come from below), we indeed have $q \in V_P(p) \Leftrightarrow (q, t) \in h_p$.

To prove the claim, we compute the height of h_p at q as the value of x_{d+1} when we plug q into the hyperplane equation, resulting in

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

As $\ell(q)$ is at height $q^\top q$, we have that the vertical distance to h_p is

$$q^\top q - 2p^\top q + p^\top p = \|q - p\|^2.$$

□

Questions

40. *What is a polytope?* Give a definition and provide a few examples.
41. *What is a face of a polytope? What is a vertex, an edge, a ridge, a facet?* Give precise definitions!
42. *Can you characterize vertex-edge graphs of 3-dimensional polytopes?* Explain Steinitz' Theorem.
43. *What is a hypercube? What is a simplex?* Define these polytope and explain what their faces are.
44. *How many* k*-faces can a* d*-dimensional polytope with* n *vertices have?* Prove a nontrivial upper bound.
45. *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.
46. *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.
47. *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.
48. *How connected is the graph of a polytope?* State and prove Balinski's theorem.
49. *What is a* d*-dimensional (Delaunay) triangulation?* Give a precise definition.
50. *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.

51. *How many facets can a* 4*-dimensional polytope with* n *vertices have?* Prove a lower bound of $\Omega(n^2)$.

52. (**This topic was not covered in this year's course in HS22 and therefore the following question will not be asked in the exam.**) *What is a* d*-dimensional Voronoi diagram?* Give a definition and explain how the Voronoi diagram relates to a polyhedron in one dimension higher!

References

[1] Michael Joswig and Thorsten Theobald, *Polyhedral and Algebraic Methods in Computational Geometry*, Universitext, Springer, 2013.

[2] Francisco Santos, A counterexample to the Hirsch conjecture. *Annals of mathematics*, pp. 383–412.

[3] Hang Si, On Monotone Sequences of Directed Flips, Triangulations of Polyhedra, and Structural Properties of a Directed Flip Graph. https://arxiv.org/abs/1809.09701, 2018.

[4] G. M. Ziegler, *Lectures on Polytopes*, vol. 152 of *Graduate Texts in Mathematics*, Springer-Verlag, Heidelberg, 1994.