Chapter 3
Polytopes and Polyhedra

Polytopes may be defined as the convex hull of finitely many points in n-dimensional space $\mathbb{R}^n$. They are fundamental objects in computational geometry. When studying polytopes, it soon becomes apparent that the proof of seemingly obvious properties often requires further clarification of the basic underlying geometric structures. An example of this is the major result that polytopes can also be represented as the intersection of finitely many affine half-spaces.

In this chapter the geometric foundations of polytopes and unbounded polyhedra will be presented from a computational viewpoint.

3.1 Definitions and Fundamental Properties

Definition 3.1 A set $P \subseteq \mathbb{R}^n$ is a *polytope* if it can be expressed as the convex hull of finitely many points. A k-dimensional polytope is called a *k-polytope*. The convex hull of $k + 1$ affinely independent points is a *k-simplex*.

A 0-dimensional polytope is just a point, a 1-polytope is a line segment and the 2-dimensional polytopes are precisely the convex polygons. We adopt the convention that the empty set is a polytope of dimension -1. See Fig. 3.1 for some examples.

From an analytical viewpoint, a polytope is a closed and bounded, and hence compact, subset of $\mathbb{R}^n$. Polytopes in lower dimensions illustrate neither the diversity of polytopes nor the depth of higher dimensional polytope theory.

We now introduce the reader to some examples of polytopes which will be useful in the following sections.

The *standard cube* C_n is the convex hull of the 2^n points which have ± 1-coordinates. If we denote the standard basis vectors in $\mathbb{R}^n$ by $e^{(1)}, \ldots, e^{(n)}$, then we can express the *cross-polytope* as the convex hull of the $2n$ points $\pm e^{(1)}, \ldots, \pm e^{(n)}$. The 3-dimensional cross-polytope is the *regular octahedron*.

The *cyclic polytopes* form an important class of polytopes with extremal properties. These properties will be discussed in further detail in Section 3.5.
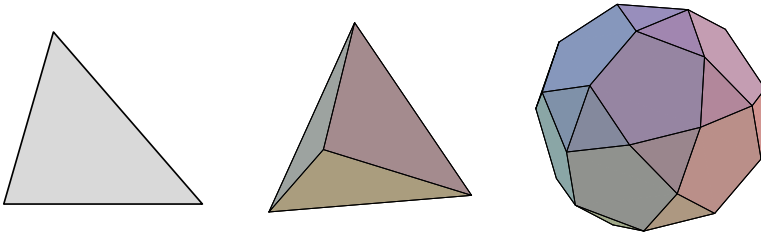
Fig. 3.1 Every 2-simplex is a triangle and every 3-simplex in $\mathbb{R}^3$ is a (generally irregular) tetrahedron. The *right hand picture* shows a 3-polytope in $\mathbb{R}^3$

Definition 3.2 The *moment curve* μ_n in $\mathbb{R}^n$ is defined as

$$\mu_n : \mathbb{R} \to \mathbb{R}^n, \quad \tau \mapsto \left(\tau, \tau^2, \ldots, \tau^n\right)^T.$$

A polytope $Z \subseteq \mathbb{R}^n$ is called *cyclic* if Z is the convex hull of points of the moment curve.

For $n = 2$ the moment curve is the standard parabola $\tau \mapsto (\tau, \tau^2)^T$. Any cyclic 2-polytope that is defined as the convex hull of $m \geq 3$ points is a convex m-gon. This is independent of the specific choice of the m points on the curve μ_2.

It is easy to verify that the image of a polytope under an affine transformation is again a polytope (of the same dimension).

Definition 3.3 An *affine automorphism* of a polytope $P \subseteq \mathbb{R}^n$ is an affine transformation of $\mathbb{R}^n$ that maps P onto itself.

The set of all affine automorphisms of a polytope is a group with respect to composition. The size of this *automorphism group* is a measure of the regularity of the polytope.

Exercise 3.4

(a) Show that for any two points $p, q \in \mathbb{R}^n$ with ± 1-coordinates there exists an affine transformation of $\mathbb{R}^n$ which leaves the standard cube fixed and maps p to q.
(b) Compute the number of affine automorphisms of the standard cube. [*Hint*: Make use of Theorem 2.7.]

3.1.1 The Faces of a Polytope

Based on the concept of supporting hyperplanes (as introduced in Section 2.3.2 or in Appendix B) we define the faces of a polytope.

Fig. 3.2 A cube in $\mathbb{R}^3$ has 8 vertices, 12 edges (which are also ridges in dimension 3) and 6 facets
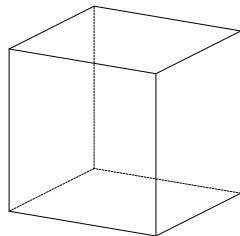
Definition 3.5 Let $P \subseteq \mathbb{R}^n$ be an n-polytope. The intersection $P \cap H$ of P with a supporting hyperplane H is called a *proper face* of P. A face of dimension k is called a *k-face*. A 0-face is called a *vertex*, a 1-face an *edge*, an $(n-2)$-face a *ridge* and an $(n-1)$-face is called a *facet*. Additionally, there are two *non-proper faces*: the empty set and P itself.

Definition 3.5 was stated for full-dimensional polytopes. The terms translate immediately to arbitrary k-polytopes $P \subseteq \mathbb{R}^n$ for $k < n$ if they are interpreted with respect to the affine hull aff P of P. An example illustrating the faces of a polytope is given in Fig. 3.2.

Theorem 3.6 *The number of faces of a polytope is finite. Faces of polytopes are polytopes themselves.*

Proof Let $P = \text{conv}\, U$ for a finite set U. For both claims it suffices to show that every proper face of P is the convex hull of a subset of U. Let H be a supporting hyperplane of P and let $U' := U \cap H$. The oriented homogeneous coordinates of H are $[a_0 : \cdots : a_n]$ and we assume without loss of generality that $P \subseteq H^+$. We will show that $H \cap P = \text{conv}\, U'$. The inclusion "$\supseteq$" is clear.

For the reverse inclusion consider a point $p = (p_1, \ldots, p_n)^T \in P \setminus \text{conv}\, U'$. There exist $u^{(1)}, \ldots, u^{(k)} \in U$ such that $p = \lambda^{(1)} u^{(1)} + \cdots + \lambda^{(k)} u^{(k)}$ with $\lambda^{(j)} \geq 0$ and $\sum \lambda^{(j)} = 1$. Here we can assume that $u^{(1)} \in U \setminus U'$ and $\lambda^{(1)} > 0$. We have to show that $p \notin H$. We have

$$a_0 + \sum_{i=1}^{n} a_i p_i = a_0 + \sum_{i=1}^{n} a_i \sum_{j=1}^{k} \lambda^{(j)} u_i^{(j)}$$

$$= a_0 + \sum_{j=1}^{k} \lambda^{(j)} \sum_{i=1}^{n} a_i u_i^{(j)} = \sum_{j=1}^{k} \lambda^{(j)} \left(a_0 + \sum_{i=1}^{n} a_i u_i^{(j)} \right),$$

where the last equation follows from $\sum_{j=1}^{k} \lambda^{(j)} = 1$. But by our assumption we have that $a_0 + \sum_{i=1}^{n} a_i u_i^{(1)} > 0$ and $a_0 + \sum_{i=1}^{n} a_i u_i^{(j)} \geq 0$ for all $j \in \{2, \ldots, n\}$. Since $\lambda^{(1)} > 0$, this implies $a_0 + \sum_{i=1}^{n} a_i p_i > 0$ or, in other words, $p \in H_\circ^+$. $\square$

Fig. 3.3 Separation of the
point $v^{(1)}$ from
conv$(V \setminus \{v^{(1)}\})$ by H and a
parallel supporting
hyperplane K
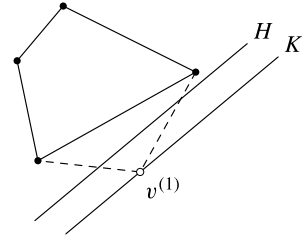
3.1.2 First Consequences of the Separating Hyperplane Theorem

As the Separating Hyperplane Theorem 2.14 is an important key to the study of the
structure of polytopes, we shall begin by studying its implications.

Theorem 3.7 *The boundary of a full-dimensional polytope $P \subseteq \mathbb{R}^n$ is the union of
all of its proper faces.*

Proof Clearly the union of all proper faces of P is contained in the boundary of P.
The reverse inclusion is implied by Corollary 2.15, which states that every boundary
point intersects at least one supporting hyperplane.   □

Theorem 3.8 *Every polytope is the convex hull of its vertices.*

Proof Let $P = \operatorname{conv} U$ for a finite set U. After successively removing all points
of U that can be expressed as a convex combination of other points in U, we obtain
a subset $V = \{v^{(1)}, \ldots, v^{(k)}\}$ that satisfies $P = \operatorname{conv} V$ and which is minimal with
respect to containment.

   We now show that every remaining point is a vertex of P. It suffices to show
this for $v^{(1)}$. Since V was chosen to be minimal, $v^{(1)}$ is not contained in the convex
hull of the other points. By Theorem 2.14 there exists an affine hyperplane H that
separates $v^{(1)}$ and $\operatorname{conv}\{v^{(2)}, \ldots, v^{(k)}\}$. We set $H = [a_0 : \cdots : a_n]$ and assume that
$v^{(1)} \in H_\circ^-$. Using the notation $\mu := a_0 + \sum_{i=1}^n a_i v_i^{(1)}$, the hyperplane K which is
parallel to H and contains $v^{(1)}$ has the oriented homogeneous coordinates $[a_0 - \mu :
a_1 : \cdots : a_n]$; see Fig. 3.3. The inequality $\mu < 0$ implies $\{v^{(2)}, \ldots, v^{(k)}\} \subseteq \operatorname{int} K^+$
and since $v^{(1)} \in K$ we have that K is a supporting hyperplane to P. Now let $p \in
P \cap K$. Since p is a convex combination of the points $v^{(j)}$, i.e., $p = \sum_{j=1}^k \lambda^{(j)} v^{(j)}$
for appropriate $\lambda^{(j)} \geq 0$ with $\sum_{j=1}^k \lambda^{(j)} = 1$, we have

$$a_0 - \mu + \sum_{i=1}^n a_i p_i = a_0 - \mu + \sum_{i=1}^n a_i \sum_{j=1}^k \lambda^{(j)} v_i^{(j)}$$

$$= \sum_{j=1}^k \lambda^{(j)} \left(a_0 - \mu + \sum_{i=1}^n a_i v_i^{(j)} \right) = 0.$$

Fig. 3.4 The point q is not contained in the affine hull of p with any face of dimension $\leq n - 2$ (here $n = 2$)
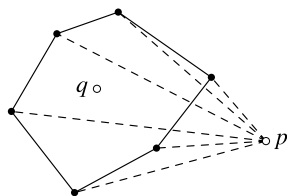
Since $\lambda^{(j)} \geq 0$ and $a_0 - \mu + \sum_{i=1}^{n} a_i v_i^{(j)} > 0$ for all $j \geq 2$, we have $\lambda^{(2)} = \cdots = \lambda^{(k)} = 0$ and also $\lambda^{(1)} = 1$. This means that $p = v^{(1)}$ and therefore that $v^{(1)}$ is a vertex of P.   □

An immediate consequence of the above theorem is that the containment-minimal set V of points that generate P is uniquely determined.

3.1.3 The Outer Description of a Polytope

The representation of a polytope as the convex hull of a finite point set is called the *$\mathcal{V}$-representation* or *inner description*. The following two central theorems state that every polytope can be equivalently described as the bounded intersection of finitely many closed half-spaces (the *$\mathcal{H}$-representation* or *outer description*). The prefixes $\mathcal{V}$- and $\mathcal{H}$- derive from the terms "vertices" and "hyperplanes".

Theorem 3.9 *Let $P \subseteq \mathbb{R}^n$ be an n-polytope, $\{F_1, \ldots, F_m\}$ the set of its facets, H_i the supporting hyperplane to P at F_i and H_i^+ the half-space containing P. Then we have*

$$P = \bigcap_{i=1}^{m} H_i^+.$$

Thus every polytope is the intersection of a finite set of closed half-spaces.

Proof The inclusion "$\subseteq$" is clear. For the inclusion "$\supseteq$" we show that every point outside of P is not contained in the intersection $\bigcap_{i=1}^{m} H_i^+$. For the following we fix a point $p \notin P$.

We study the set $\{G_1, \ldots, G_k\}$ of all faces of P of dimension $\leq n - 2$. Let q be a point in the interior of P which is not contained in the set $\bigcup_{i=1}^{k} \mathrm{aff}(G_i \cup \{p\})$. Such a point exists since the interior of an n-polytope has dimension n and can therefore not be covered by a finite number of affine subspaces of dimension $\leq n - 1$ (see Fig. 3.4). The segment $[p, q]$ intersects the boundary of P in a uniquely determined point z which, by Theorem 3.7, is contained in a proper face of P. By the choice of q it is guaranteed that z is not contained in a face of dimension $j < n - 1$. This implies that there exists an $i \in \{1, \ldots, m\}$ with $z \in F_i$. So we have $z \in H_i$ and $q \in H_i^+$, but $p \in H_i^- \setminus H_i$, i.e., $p \notin \bigcap_{i=1}^{m} H_i^+$.   □

When, as in Theorem 3.9, the polytope $P \subseteq \mathbb{R}^n$ is full-dimensional, the affine span of every facet F defines a hyperplane H. Assuming that H has the form $H = [a_0 : \cdots : a_n]$ and $P \subseteq H^+$, every positive multiple of $(a_1, \ldots, a_n)^T$ is called an *inner normal vector* of F and every negative multiple of $(a_1, \ldots, a_n)^T$ is called an *outer normal vector* of F. If dim $P < n$, then for any facet F, there exist infinitely many affine hyperplanes of $\mathbb{R}^n$ that contain F.

We can now refine Theorem 3.7, which states that the boundary of a polytope is the union of its facets.

Theorem 3.10 *If the intersection P of a finite number of closed affine half-spaces in $\mathbb{R}^n$ is bounded, then P is a polytope.*

Proof The proof is completed by induction over the dimension n of the space. The statement is clear for dimension ≤ 1 . So let $n \geq 2$ and

$$P = \bigcap_{i=1}^{m} H_i^+$$

be the bounded intersection of a finite number of affine half-spaces in $\mathbb{R}^n$. Let $F_j := H_j \cap P$, $j \in \{1, \ldots, m\}$. Then F_j is a bounded intersection of half-spaces in the hyperplane H_j. Since H_j can be identified with an affine space of dimension $n - 1$, we know by the inductive hypothesis that F_j is a polytope in H_j and therefore also a polytope in $\mathbb{R}^n$. Let V_j be the set of vertices of F_j and $V = \bigcup_{j=1}^{m} V_j$.

It suffices to show that $P = \text{conv } V$. The inclusion "$\supseteq$" is clear since $V \subseteq P$ and P is convex. For the reverse inclusion consider a point $q \in P$. If q is a boundary point of P, then there exists a $j \in \{1, \ldots, m\}$ with $q \in F_j$. The point q is therefore a convex combination of V_j which in particular implies that $q \in \text{conv } V$. If q is contained in the interior of P, then q is contained in a segment, $[r, s]$, formed by the intersection of a line with P. Since r and s are on the boundary of P they are contained in conv V and thus $q \in \text{conv } V$. $\square$

Example 3.11 The hyperplanes that define facets of the standard cube C_n are precisely $H_i = [1 : h_1^{(i)} : \cdots : h_n^{(i)}]$ for $i \in \{1, \ldots, 2n\}$ with

$$h_k^{(i)} = \begin{cases} 1 & \text{if } i = k, \\ -1 & \text{if } i = k + n, \\ 0 & \text{otherwise} \end{cases}$$

for $k \in \{1, \ldots, n\}$.

Computing the $\mathcal{H}$-representation when given the $\mathcal{V}$-representation of a polytope and vice versa is a major topic of computational geometry and will be discussed in Chapter 5.

Exercise 3.12 Show that the intersection of a polytope with an arbitrary affine sub-space is a polytope.

Exercise 3.13 For a polytope P show that:

(a) The intersection of a set of faces of P is a face of P.
(b) Every ridge of P is the intersection of exactly two facets of P.
(c) If G is a face of P and F is a face of G, then F is a face of P.

3.2 The Face Lattice of a Polytope

Containment defines a partial order on the set $\mathcal{F}(P)$ of all faces of a polytope P. Theorem 3.6 tells us that this set is finite, i.e., $(\mathcal{F}(P), \subseteq)$ is a finite *partially ordered set* (or *poset*). As a purely combinatorial object this poset is an important interface between the analytically focused general complexity theory and discrete geometry.

Exercise 3.14 Show that $(\mathcal{F}(P), \subseteq)$ satisfies the following conditions:

(a) There exists a uniquely determined smallest and largest face of P.
(b) For two arbitrary faces $F, G \in \mathcal{F}(P)$ there exists a uniquely determined small-est face $F \vee G$ such that $F \subseteq F \vee G$ and $G \subseteq F \vee G$.
(c) For two arbitrary faces $F, G \in \mathcal{F}(P)$ there exists a uniquely determined largest face $F \wedge G$ such that $F \supseteq F \wedge G$ and $G \supseteq F \wedge G$.

The properties described in Exercise 3.14 show that $(\mathcal{F}(P), \subseteq)$ is a *lattice*, called the *face lattice* of P.

Definition 3.15 A *combinatorial isomorphism* of two polytopes is a (poset-)iso-morphism of the face lattices. If there exists such a combinatorial isomorphism, we call the two polytopes *combinatorially equivalent*. The *combinatorial type* of a polytope is the isomorphism type of its face lattice.

Exercise 3.16 Show that every affine transformation of a polytope P to a poly-tope Q induces an isomorphism from $\mathcal{F}(P)$ to $\mathcal{F}(Q)$.

Exercise 3.17 Give an example of two combinatorially equivalent polytopes such that there does not exist an affine transformation that maps one to the other.

Theorem 3.18 *Let F and G be faces of P such that $F \subseteq G$. Then*

$$\mathcal{F}(F, G) := \left\{ F' \in \mathcal{F}(P) : F \subseteq F' \subseteq G \right\}$$

with the partial order induced by containment, is isomorphic to the face lattice of a polytope of dimension $\dim G - \dim F - 1$.

Fig. 3.5 Two examples of polytopes $P \cap H(\epsilon)$ for $\dim F \in \{0, 1\}$. (Notation as in the proof of Theorem 3.18.) The polytope P is a *bipyramid* over a pentagon
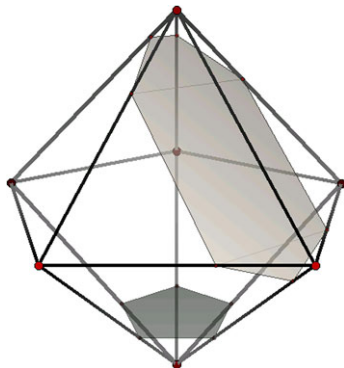
Proof Since Theorem 3.6 established that every face of a polytope is a polytope itself, we can assume without loss of generality that $G = P$. Let P be full-dimensional. We choose F as a proper face since otherwise there is nothing left to show.

Let V be the vertex set of P and $V(F) = F \cap V$ the vertex set of the face F. Choose a supporting hyperplane H to P with $F = P \cap H$. We assume that H has the oriented homogeneous coordinates $[a_0 : \cdots : a_n]$ and that $P \subseteq H^+$ holds. For every sufficiently small $\epsilon > 0$ we have that the hyperplane $H(\epsilon) = [a_0 - \epsilon : a_1 : \cdots : a_n]$, which is parallel to H, separates the vertex set $V(F)$ from its complement: $V(F) \subseteq \operatorname{int} H(\epsilon)^-$ and $V \setminus V(F) \subseteq \operatorname{int} H(\epsilon)^+$. See Fig. 3.5.

Let x be a point in the relative interior of F. The hyperplane $H(\epsilon)$ contains an interior point, y say, of P. Now let A be an $(n - \dim F)$-dimensional affine subspace containing x and y but no point in $\operatorname{aff} F$ other than x. That is, $\operatorname{aff} F$ and A are complementary affine subspaces meeting at x. Then $A \cap H(\epsilon)$ is an affine subspace of dimension $n - \dim F - 1$ which is affinely generated by the set

$$P(F, A, \epsilon) := P \cap A \cap H(\epsilon),$$

which by Theorem 3.10 is a polytope. The map

$$\alpha : \mathcal{F}(F, P) \to \mathcal{F}\big(P(F, A, \epsilon)\big) : F' \mapsto F' \cap A \cap H(\epsilon)$$

respects containment and is bijective since $\alpha^{-1}(F' \cap A \cap H(\epsilon)) = \operatorname{aff}((F' \cap A \cap H(\epsilon)) \cup F) \cap P = \operatorname{aff}(F') \cap P = F'$. Since $\mathcal{F}(F, P)$ does not depend on x, y, A or ϵ we have that the combinatorial type of the polytope $P(F, A, \epsilon)$ is independent of x, y, A and ϵ. $\qquad\square$

Given a face F of P we call the polytope $P(F, A, \epsilon)$ a *face figure* of F. An implication of Theorem 3.18 is the fact that $\mathcal{F}(F, G)$, for $\dim G - \dim F = 2$, is always the face lattice of a segment, i.e., there exist exactly two faces E_1 and E_2 of dimension $\dim F + 1$ which lie between F and G. This property is called the *diamond property* of $\mathcal{F}(P)$; see Fig. 3.6. For a bipyramid over a pentagon, Fig. 3.5 shows the (pentagonal) face figure of a vertex, also called a *vertex figure*; the shaded

Fig. 3.6 The diamond
property of the face lattice.
The faces E_1, E_2, F, G
satisfy $E_1 \vee E_2 = G$ and
$E_1 \wedge E_2 = F$
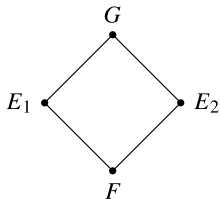
heptagon corresponds to the face figure of an edge, which is a segment, as discussed
above.

Let $P \subseteq \mathbb{R}^n$ be an n-polytope and let $f_k(P)$ be the number of k-dimensional faces
of P. Then $f(P) := (f_0(P), f_1(P), \ldots, f_{n-1}(P))$ is called the f-*vector* of P. The
f-vector is clearly a combinatorial invariant since it only depends on the combina-
torial type of P. An interesting—and very complicated—task is to determine which
n-tuples of natural numbers can be f-vectors of n-polytopes.

Exercise 3.19 Compute the f-vector of the n-dimensional standard cube C_n and
describe its face lattice.

One may ask what "typical" polytopes look like. A more rigorous statement of
this naive question can be formulated in several ways using stochastic terms. As
an example, we will study the convex hulls of random points on the unit sphere in
Section 3.6. In many cases the term "typical" corresponds to "general position".

Exercise 3.20 Let $K \subseteq \mathbb{R}^n$ be a full-dimensional convex set. Show that a finite set
X of uniformly distributed random points from K is almost certainly in *general
position*, i.e. the probability of $n + 1$ of these points being affinely independent is 1.
In particular this implies that every proper face of conv X is a simplex.

The last property inspires the following definition.

Definition 3.21 A polytope P is called *simplicial* if all proper faces of P are sim-
plices. It is called *simple* if the face figure of every proper face of P is a simplex.

The cross polytopes conv$\{\pm e^{(1)}, \ldots, \pm e^{(n)}\}$ are simplicial, while the cubes C_n
are simple. The relationship between these two properties, simplicial and simple,
will be clarified in Section 3.3.

Exercise 3.22 Show that a polytope is both simplicial and simple if and only if it is
a simplex or a polygon.

Exercise 3.23 Let P be an n-polytope with vertex set V and edge set E. The *graph*
$\Gamma(P)$ is the abstract graph (V, E) with natural incidence. Show:

(a) The graph $\Gamma(P)$ is connected.
(b) Every vertex is incident with at least n edges.

(c) The n-polytope P is simple if and only if every vertex is incident with exactly n edges.

3.3 Polarity and Duality

In the following section we introduce the concept of polarity. Given a polytope P which contains the origin in its interior, we assign to P a polar polytope P° such that every k-face of P corresponds to an $(n-k-1)$-face of P°. In particular, we have that $f_{n-i-1}(P) = f_i(P^\circ)$.

Example 3.24 For the standard cube $C_3 = [-1, 1]^3$ in $\mathbb{R}^3$ we have $f_0(C_3) = 8$, $f_1(C_3) = 12$, $f_2(C_3) = 6$. For the three-dimensional cross-polytope (the octahedron) $Q = \mathrm{conv}\{\pm e^{(i)} : 1 \le i \le 3\}$ (where $e^{(i)}$ denotes the i-th standard basis vector), we have $f_0(Q) = 6$, $f_1(Q) = 12$, $f_2(Q) = 8$ (see Fig. 3.7).
    We will see in Example 3.30 that Q is the polar polytope of C_3.

As in Section 2.3.1, let $\langle \cdot, \cdot \rangle$ denote the Euclidean scalar product and let $\| \cdot \|$ with $\|x\| := \langle x, x \rangle^{1/2}$ be the Euclidean norm.

Definition 3.25 For $X \subseteq \mathbb{R}^n$ the *polar set* X° is defined as

$$X^\circ = \{y \in \mathbb{R}^n : \langle x, y \rangle \le 1 \text{ for all } x \in X\}.$$

Exercise 3.26 Show that $X \subseteq Y$ implies $Y^\circ \subseteq X^\circ$ for $X, Y \subseteq \mathbb{R}^n$.

Proposition 3.27 *Let $X \subseteq \mathbb{R}^n$. Then X° is closed and convex and $0 \in X^\circ$.*

Proof Clearly $0 \in X^\circ$. Let $x \in \mathbb{R}^n \setminus \{0\}$, then

$$\{x\}^\circ = \{y \in \mathbb{R}^n : \langle x, y \rangle \le 1\} = [1 : -x_1 : \cdots : -x_n]^+$$

is a closed affine half-space and $\{0\}^\circ = \mathbb{R}^n$. The intersection $X^\circ = \bigcap_{x \in X} \{x\}^\circ$ of closed and convex sets is again closed and convex. $\square$
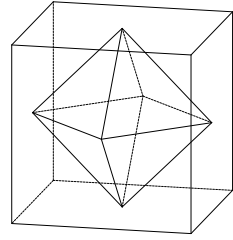
Theorem 3.28 *If $P \subseteq \mathbb{R}^n$ is an n-polytope with $0 \in \mathrm{int}\, P$, then P° is also an n-polytope with $0 \in \mathrm{int}\, P$. We have*

$$P^\circ = \bigcap_{v \in V} \{y \in \mathbb{R}^n : \langle v, y \rangle \le 1\} = \bigcap_{v \in V} [1 : -v_1 : \cdots : -v_n]^+, \qquad (3.1)$$

where V is the vertex set of P.

Proof Since P is bounded, we have that P is contained in an open ball $B(0, \rho)$ with center 0 and radius ρ. For all $x \in \mathbb{R}^n$ with $\|x\| \le 1/\rho$ the Cauchy–Schwarz

Fig. 3.7 The cube $[-1, 1]^3$
and octahedron
conv$\{e^{(i)} : 1 \leq i \leq 3\}$

inequality gives

$$\left|\langle x, y \rangle\right| \leq \|x\| \, \|y\| \leq \frac{1}{\rho} \|y\| \leq 1 \quad \text{for all } y \in P,$$

and thus the ball $B(0, 1/\rho)$ is contained in P°. As a consequence, P° is full-dimensional. Since P contains a ball $B(0, \rho')$ we can analogously deduce that P° is bounded.

Equation (3.1) remains to be proven. The inclusion "$\subseteq$" follows immediately from Definition 3.25. For the reverse inclusion "$\supseteq$" consider a point y that is not contained in P°. An arbitrary point $x \in P$ can be expressed as a convex combination $\sum_{i=1}^{k} \lambda^{(i)} v^{(i)}$ of vertices of P. Clearly we have

$$\langle x, y \rangle = \sum_{i=1}^{k} \lambda^{(i)} \langle v^{(i)}, y \rangle \leq \max\{\langle v^{(i)}, y \rangle : 1 \leq i \leq k\},$$

where the last inequality follows from $\sum_{i=1}^{k} \lambda^{(i)} = 1$. If $\langle x, y \rangle > 1$ then there exists a vertex $v^{(i)}$ such that $\langle v^{(i)}, y \rangle > 1$, which proves the statement. $\square$

Theorem 3.29 *For an n-polytope $P \subseteq \mathbb{R}^n$ with $0 \in \text{int } P$ we have:*

(a) $(P^\circ)^\circ = P$.
(b) *For every boundary point p of P the affine hyperplane*

$$H = \left\{ x \in \mathbb{R}^n : \langle p, x \rangle = 1 \right\}$$

is a supporting hyperplane to P°.

Proof (a) The definition of polarity implies immediately that $P \subseteq (P^\circ)^\circ$. For the reverse direction let $P = \bigcap_{i=1}^{m} H_i^+$ and let x be a point not contained in P. Then there exists an $i \in \{1, \ldots, m\}$ with $x \notin H_i^+$. By the Separation Theorem 2.14 there exists a $v \in \mathbb{R}^n$ with $\langle v, x \rangle > 1$. Since $\langle v, y \rangle \leq 1$ for all $y \in H_i^+$, we have $v \in P^\circ$ and since $\langle v, x \rangle > 1$ we have $x \notin (P^\circ)^\circ$.

(b) For every $p \in P \setminus \{0\}$, $H^+ = [1 : -p_1 : \cdots : -p_n]^+$ is a half-space which contains the polytope P°. If p is a boundary point of P then, by Theorem 3.7, it belongs to a face of P and there exists a vector $x \in \mathbb{R}^n$ such that the hyperplane $H' = \{y \in \mathbb{R}^n : \langle x, y \rangle = 1\}$ supports P and contains the point p. We now have

$x \in P^\circ$ and $x \in H$ such that H intersects the polytope P° and thus H is a supporting hyperplane to P°. □

Example 3.30 The description of the hyperplanes that define the facets of the standard cube C_n in Example 3.11 shows that C_n is polar to the n-dimensional cross-polytope.

Lemma 3.31 *Let $P \subseteq \mathbb{R}^n$ be an n-polytope with $0 \in \operatorname{int} P$. For every proper face F of P the face*

$$F^* := \left\{ x \in P^\circ : \langle x, y \rangle = 1 \text{ for all } y \in F \right\} \subsetneq F^\circ$$

is a proper face of P°.

Proof For every $p \in F$, Theorem 3.29b implies that the hyperplane $H = \{x \in \mathbb{R}^n : \langle p, x \rangle = 1\}$ is a supporting hyperplane of P°, and thus $P^\circ \cap H$ is a face of P°. Since F is the convex hull of a finite number of points, we can express F^* as an intersection of a finite number of such faces of P°. Hence F^* is a face of P°. By construction the face F^* is contained in the set F°. □

Lemma 3.31 induces a map $\phi : F \mapsto F^*$ from the set of all proper faces of P to the set of all proper faces of P°, with $\phi(P) = \emptyset$ and $\phi(\emptyset) = P^\circ$.

Theorem 3.32 *Let $P \subseteq \mathbb{R}^n$ be an n-polytope with $0 \in \operatorname{int} P$. The map ϕ is bijective and for all $k \in \{0, \dots, n-1\}$ it maps the k-faces of P to the $(n-k-1)$-faces of P°. Furthermore ϕ is containment-reversing, i.e., $F \subseteq G$ implies $G^* \subseteq F^*$.*

Proof Exercise 3.26 implies that ϕ is containment-reversing.

Since the map ϕ can also be applied to the faces of P°, in order to prove bijectivity we know by Theorem 3.29a that it suffices to show that $\phi(\phi(F)) = F$ for every face F of P. For the non-proper faces this is satisfied by definition. For every proper face F we have by definition that

$$\phi\big(\phi(F)\big) = \left\{ x \in \mathbb{R}^n : \langle x, y \rangle = 1 \text{ for all } y \in \phi(F) \right\}$$

with

$$\phi(F) = \left\{ y \in P^\circ : \langle x, y \rangle = 1 \text{ for all } x \in F \right\},$$

and thus $F \subseteq \phi(\phi(F))$.

For the reverse inclusion, consider a point $p \in P$ with $p \notin F$. If we denote by $H = [1 : -h_1 : \cdots : -h_n]$ a supporting hyperplane to P that contains F, then $p \in H_o^+$ and hence $\langle p, h \rangle < 1$. Notice the first coordinate of H equals 1 (up to scaling by a positive real number) as the origin is an interior point of P. Since $h \in \phi(F)$ it then follows that $p \notin \phi(\phi(F))$.

Our dimension statement remains to be proven. For non-proper faces it is clearly satisfied. Every proper k-face F contains $k + 1$ affinely independent points such

Fig. 3.8 A polytope
$P = \mathrm{conv}\{p^{(1)}, \ldots, p^{(4)}\}$ and
its polar polytope
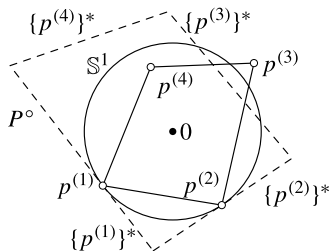
that $\phi(F)$ is contained in the intersection of $k + 1$ hyperplanes whose equations are linearly independent. This implies $\dim \phi(F) \leq n - k - 1$, but since $F \subseteq \phi(\phi(F))$, this must be an equality. $\square$

A bijection that reverses the order relation of a lattice or poset is called an *anti-isomorphism*.

Corollary 3.33 *Let* $P \subseteq \mathbb{R}^n$ *be an n-polytope with* $0 \in \mathrm{int}\, P$. *The polytope* P *is simplicial if and only if the polar polytope* P° *is simple.*

Example 3.34 Let $P = \mathrm{conv}\{p^{(1)}, \ldots, p^{(4)}\} \subseteq \mathbb{R}^2$ be the quadrangle depicted in Fig. 3.8. Then the polar polytope $P^\circ = \bigcap_{i=1}^{4}\{x \in \mathbb{R}^2 : \langle p^{(i)}, x \rangle \leq 1\}$ is the dashed quadrangle. The segment $[0, p^{(i)}]$ is perpendicular to the line

$$H_i := \{x \in \mathbb{R}^2 : \langle p^{(i)}, x \rangle = 1\} = \mathrm{aff}\big(\{p^{(i)}\}^*\big).$$

For $i \in \{1, 2\}$, $p^{(i)}$ lies on the unit circle $\mathbb{S}^1$ such that the line H_i is tangent to the unit circle. The point $p^{(3)}$ lies outside of the unit circle such that H_3 intersects the interior of the unit circle. The fourth point $p^{(4)}$ lies in the interior of the unit circle such that H_4 lies completely outside of the unit circle. The distance from the origin to the line H_i is always the reciprocal of the distance between 0 and $p^{(i)}$.

In this section we often assumed that the polytope P contained the origin as an interior point. By restricting to $\mathrm{aff}\, P$, and via a suitable translation, we can without any loss of generality always assume this to be the case, i.e., every polytope has an affine image which satisfies this condition. This implies that for every polytope P there exists a polytope P' whose face lattice $\mathcal{F}(P')$ is anti-isomorphic to $\mathcal{F}(P)$. Such a polytope is said to be *dual* to P.

3.4 Polyhedra

Polytopes are the elementary building blocks of computational geometry, but it is often more natural to study a wider class of objects: polyhedra. This will be of particular importance in Chapter 4, as well as other chapters, when we discuss linear programming.

Definition 3.35 A set $P \subseteq \mathbb{R}^n$ is called a *polyhedron* if it can be represented as the intersection of a finite number of closed affine half-spaces.

Thus, a polytope is a bounded polyhedron. In general we cannot describe a polyhedron as the convex hull of a finite number of points. The most basic example of this is a single half-space. Nevertheless, the differences between polytopes and unbounded polyhedra are manageable.

To do this we distinguish between two kinds of unbounded polyhedra: A polyhedron P either contains an affine line or it does not. In the latter case we call P *pointed*. First assume that $P = H_1^+ \cap \cdots \cap H_k^+$ is pointed. Then $n \leq k$ and $H_1 \cap \cdots \cap H_k$ is either empty or contains exactly one point. Without loss of generality we may assume that the first n hyperplanes intersect in a point, z. If H_i^+ has homogeneous coordinates $[h_0^{(i)} : \cdots : h_n^{(i)}]^+$ this means that

$$z = \left[h_0^{(1)} : \cdots : h_n^{(1)}\right] \cap \cdots \cap \left[h_0^{(n)} : \cdots : h_n^{(n)}\right].$$

The point z may or may not be contained in P. To get a clearer image of P we apply an affine transformation to P which transforms each hyperplane H_i to the coordinate hyperplane $E_i = [0 : \cdots : 0 : 1 : 0 : \cdots : 0] = \{x \in \mathbb{R}^n : x_i = 0\}$ for $1 \leq i \leq n$. In this way, z will automatically be mapped to the origin. The affine transformation described above can be represented most conveniently by the $(n+1) \times (n+1)$-matrix

$$T = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ h_0^{(1)} & h_1^{(1)} & \cdots & h_n^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ h_0^{(n)} & h_1^{(n)} & \cdots & h_n^{(n)} \end{pmatrix},$$

which operates on the left, as usual. That T indeed has all the properties required is a consequence of Exercise 2.17. By our choice of coordinates, $E_1, \ldots, E_n$ are oriented so that $E_1^+ \cap \cdots \cap E_n^+$ is the positive orthant. Hence the transformed polyhedron $[T]P$ is contained in the positive orthant. Now consider a further projective transformation, defined by the non-negative matrix

$$B = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 & 1 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & & \ddots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}.$$

The map $[B]$ is not an affine transformation. It maps the ideal hyperplane $[1 : 0 : \cdots : 0]$ to the projective hyperplane $[1 : -1 : \cdots : -1]$ so that the coordinate hyperplanes stay fixed. Furthermore the image of the positive orthant under the map $[B]$ is the n-simplex

$$E_1^+ \cap \cdots \cap E_n^+ \cap [1 : -1 : \cdots : -1]^+.$$

In particular, the image $[BT]P$ is a bounded polyhedron, i.e., a polytope. We have now proved the following theorem.

Theorem 3.36 *Every pointed polyhedron is projectively equivalent to a polytope.*

Pointed polyhedra can be imagined as polytopes with a specific proper face that has been moved to the ideal hyperplane. For an example computation, see Section 3.6.3 below.

Note that the image of a polyhedron under a projective transformation is not necessarily a polyhedron. However, the following case, which is most relevant for us, does not have this problem.

Exercise 3.37 Let $P \subseteq \mathbb{R}_{\geq 0}^n$ be a polyhedron in the positive orthant and $[A]$ the projective transformation to a matrix $A \in \mathrm{GL}_{n+1} \mathbb{R}$ with non-negative coefficients. Show that the image $[A]P$ is again a polyhedron.

We still need to consider the case where P is not pointed. In this case we choose an affine subspace A of $\mathbb{R}^n$ which is contained in P and which is maximal with respect to dimension. The linear subspace L of $\mathbb{R}^n$ which is parallel to A is called the *lineality space* of P. Let p be an arbitrary point of P and A' be the affine orthogonal complement of A that contains p. The intersection $P \cap A'$ is a polyhedron which contains no affine line and is therefore pointed.

Definition 3.38 For $X, Y \subseteq \mathbb{R}^n$, the *Minkowski sum* of X and Y is defined as

$$x + y = \{x + y : x \in X,\ y \in Y\}.$$

The Minkowski sum is called *direct* if $x + y = v + w$ with $x, v \in X$ and $y, w \in Y$ implies $x = v$ and $y = w$.

Using this notation, the lineality space, as defined above, gives us a direct Minkowski sum $P = (P \cap A') + L$, establishing the following lemma.

Lemma 3.39 *Every polyhedron can be expressed as the direct Minkowski sum of a pointed polyhedron and a linear subspace.*

In this decomposition it is possible that the pointed polyhedron or the lineality space is just a single point. In those cases the decomposition as a Minkowski sum is trivial.

In summary, we can say that statements about polyhedra can be traced back to statements about polytopes. As an example of this consider the generalization of Theorem 3.8 which is discussed in Exercise 3.41. However, first consider two further definitions.

Definition 3.40 Let $A \subseteq \mathbb{R}^n$. A *positive combination* of A is a linear combination $\sum_{i=1}^{m} \lambda^{(i)} a^{(i)}$ with $a^{(i)} \in A$ and $\lambda^{(i)} \geq 0$ for all i. The set of all positive combinations of A is called the *positive hull* of A, which we denote by $\mathrm{pos}\,A$.

Fig. 3.9 The positive hull of
a finite point set
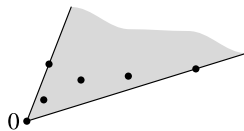
The positive hull of a set A is a *convex cone* in the sense that pos A is convex and
that $a + b$ and λa are contained in pos A for $a, b \in$ pos A and $\lambda \geq 0$ (see Fig. 3.9).
In a convex cone K we call a half-line $x + \mathbb{R}_{\geq 0} y \subseteq \partial K$ with $x \in K$ and $y \in \mathbb{R}^n \setminus \{0\}$
a *ray* of K.

Exercise 3.41 Every polyhedron $P \subseteq \mathbb{R}^n$ can be expressed as a Minkowski sum

$$P = \text{conv } V + \text{pos } R$$

for finite sets V and R.

Exercise 3.42 Show that the cone pos R in Exercise 3.41 is uniquely determined. Is
the polytope conv V also unique, in general?

The cone pos R in the preceding exercises is called the *recession cone* of the
polyhedron P.

Exercise 3.43 The *product*

$$\left\{ (p, q) \in \mathbb{R}^{n+n'} : p \in P, \, q \in Q \right\}$$

of two polyhedra $P \subseteq \mathbb{R}^n$ and $Q \subseteq \mathbb{R}^{n'}$ is a polyhedron.

3.5 The Combinatorics of Polytopes

As mentioned in Chapter 1, computational problems often require the transforma-
tion from a $\mathcal{V}$-representation to an $\mathcal{H}$-representation and vice versa. Before we study
explicit algorithms for this task in Chapter 5, it is necessary to improve our under-
standing of the combinatorial structure of polytopes.

To be able to discuss the complexity of an algorithm it is necessary to first deter-
mine how large the output of an algorithm may be in relation to its input. For convex-
hull-algorithms, i.e., methods to compute the facets of a convex hull of a given point
set, we have to answer the question of how many facets an n-dimensional polytope
with m vertices may have. The reverse question, how many vertices an n-polytope
with m facets can have, is equivalent by polarization. As before, let $f_k(P)$ denote
the number of k-dimensional faces of an n-polytope P for $-1 \leq k \leq n$. In particular
we have $f_{-1}(P) = f_n(P) = 1$.

The Upper-bound Theorem, a fundamental result in polytope theory, states that
the cyclic polytopes from Definition 3.2 are extremal in the following sense. Let

$Z_n(m)$ denote a cyclic polytope in $\mathbb{R}^n$ formed by the convex hull of m points on the moment curve. This notation purposefully neglects which m points define the cyclic polytope. At the end of this section this simplification will be justified (in Exercises 3.49–3.51) by the statement that two such cyclic polytopes are combinatorially equivalent.

Theorem 3.44 (Upper-bound Theorem, McMullen 1970) *An n-dimensional polytope with m vertices has at most as many k-faces as a cyclic polytope $Z_n(m)$ for all $k \in \{-1, \ldots, n\}$.*

In Exercises 3.49–3.51 we will compute the number of facets of cyclic polytopes. This yields the following explicit upper bound.

Corollary 3.45 *The number of facets of an n-dimensional polytope with m vertices is bounded by*

$$
\begin{cases}
\dfrac{m}{m - \frac{n}{2}} \binom{m - \frac{n}{2}}{m - n} & \text{if } n \text{ is even}, \\[2ex]
2 \binom{m - \frac{n+1}{2}}{m - n} & \text{if } n \text{ is odd}.
\end{cases}
$$

When we switch to the dual picture we obtain the same upper bound for the number of vertices of an n-polytope with m facets.

We will not fully prove Theorem 3.44 in this section. Instead we give a proof for an upper bound which has the right order of magnitude for the number of facets.

Theorem 3.46 *An n-polytope with m vertices has at most $2\binom{m}{\lfloor n/2 \rfloor}$ facets and in total not more than $2^{n+1}\binom{m}{\lfloor n/2 \rfloor}$ faces. For fixed n, both numbers have the same order of magnitude $O(m^{\lfloor n/2 \rfloor})$.*

We will prove this statement first for simplicial polytopes and then we deduce the non-simplicial case as a corollary.

Lemma 3.47 *For a simplicial n-polytope P we have*:

(a) $(n - k) f_k(P) \leq \binom{n}{k+1} f_{n-1}(P)$ *for* $k \in \{-1, \ldots, n\}$;
(b) $n f_0(P) + (n - 1) f_1(P) + \cdots + 2 f_{n-2}(P) \leq (2^n - 2) f_{n-1}(P)$;
(c) $f_{n-1}(P) \leq 2 f_{\lfloor n/2 \rfloor - 1}(P)$.

Proof For the first statement we count the number of k-faces that are incident to a given facet of P and vice versa. By our assumption every facet is an $(n-1)$-simplex which contains exactly $\binom{n}{k+1}$ k-faces. On the other hand we have that the face figure of a k-face is an $(n - k - 1)$-polytope that has at least $n - k$ facets. This implies the first statement. The second statement follows from the first by summation over k from 0 to $n - 2$.

For the third statement, we consider the dual polytope P' which is by Corollary 3.33 simple. We have to show that $f_0(P') \leq 2 f_{\lceil n/2 \rceil}(P')$.

Now we will limit the number of vertices of P' with respect to the number of $\lceil n/2 \rceil$-faces. After an affine transformation we can assume without loss of generality that no two edges of P' have the same x_n-coordinate. In the following we imagine that the n-coordinate is "pointing upwards".

Consider a vertex v and the n edges incident with v. Then there are at least $\lceil n/2 \rceil$ edges that point downwards or at least $\lceil n/2 \rceil$ edges that point upwards. In the first case we have that every $\lceil n/2 \rceil$-tuple of upward pointing edges determines a $\lceil n/2 \rceil$-face for which v is the lowest vertex. In the second case we have that each $\lceil n/2 \rceil$-tuple of downward pointing edges determines a $\lceil n/2 \rceil$-face for which v is the highest vertex. Since the lowest and highest vertex for each face are unique, there are at most twice as many vertices as there are $\lceil n/2 \rceil$-faces.   □

Lemma 3.48 *For each n-polytope P there exists an n-dimensional simplicial polytope Q with the same number of vertices as P such that $f_k(Q) \geq f_k(P)$ for $1 \leq k \leq n$.*

Proof We can assume that $P \subseteq \mathbb{R}^n$. Our goal is to obtain the polytope Q from P by slightly moving all the vertices.

For the *perturbation* of one vertex v we employ the following operation. Pick an affine hyperplane H which strictly separates v from all the other vertices of P. We may orient H so that $H_\circ^-$ contains v. Now choose a point $v' \in \operatorname{int} P \cap H_\circ^-$ which is not contained in a hyperplane spanned by any $n+1$ vertices of P. Replacing v by v' we obtain the polytope

$$P' = \operatorname{conv}\big(\{w : w \text{ vertex of } P \text{ distinct from } v\} \cup \{v'\}\big),$$

which is contained in P. We want to show that P' has at least as many faces of each dimension as P. To this end we will describe an injective map ι from the faces of P to the faces of P'.

Let F be a proper face of P and let A be an affine hyperplane supporting P with $A \cap P = F$. Notice that, as P contains P', the hyperplane A does not separate P'. If F does not contain the vertex v then we set $\iota(F) = A \cap P' = A \cap P = F$, and this is a face of P'.

It remains to consider the case when F contains v. If F is a simplex, then

$$\iota(F) = \operatorname{conv}\big(\{w : w \text{ vertex of } F \text{ distinct from } v\} \cup \{v'\}\big)$$

is a face of P'. We may then assume that F is not a simplex. As v' is not contained in F in this case it follows that $\iota(F) = A \cap P'$ is a face of P' of the same dimension as F. This yields a dimension-preserving map ι from the face lattice of P to the face lattice of P'. It is easy to see that ι is injective.

To construct the polytope Q we pick a linear ordering $v^{(1)}, \ldots, v^{(m)}$ of the vertices of P. Inductively perturbing the vertices in this order gives a sequence of n-dimensional polytopes $P^{(1)}, \ldots, P^{(m)}$ all of which have precisely m vertices. Setting $P^{(0)} = P$ and $Q = P^{(m)}$ we have

$$f_k(P) \leq f_k\big(P^{(i)}\big) \leq f_k(Q)$$

for $1 \leq i \leq m$ and $1 \leq k \leq n$. Moreover, our procedure guarantees that the vertices of Q are in general position and therefore Q is simplicial. □

Proof of Theorem 3.46 By Lemma 3.48 it suffices to consider simplicial polytopes P. Since the number of $(\lfloor n/2 \rfloor - 1)$-faces clearly satisfies

$$f_{\lfloor n/2 \rfloor - 1}(P) \leq \binom{m}{\lfloor n/2 \rfloor},$$

Lemma 3.47c implies

$$f_{n-1}(P) \leq 2\binom{m}{\lfloor n/2 \rfloor}$$

and using Lemma 3.47b we obtain

$$f_0(P) + f_1(P) + \cdots + f_n(P) \leq 2^{n+1}\binom{m}{\lfloor n/2 \rfloor}. \qquad \square$$

To conclude this section we will study the cyclic polytopes introduced in Definition 3.2. As mentioned in Theorem 3.44, these polytopes maximize the f-vector of all polytopes.

Exercise 3.49 Show that each set of n points on the moment curve in $\mathbb{R}^n$ are affinely independent. This implies that cyclic polytopes are simplicial.

As a result of the following exercise (and Exercise 3.55) we know that two cyclic polytopes of the same dimension and the same number of vertices are combinatorially equivalent. This justifies the notation $Z_n(m)$.

Exercise 3.50 (Gale Evenness Condition)   Let V be the vertex set of a cyclic polytope in $\mathbb{R}^n$ with the induced order $\prec$ with respect to the moment curve, i.e., $x(\tau_1) \prec x(\tau_2)$ if and only if $\tau_1 < \tau_2$. Let $U = \{v^{(1)}, \ldots, v^{(n)}\} \subseteq V$ be an n-tuple of vertices of P, where $v^{(1)} \prec v^{(2)} \prec \cdots \prec v^{(n)}$. Show that conv U is a facet of P if and only if for every two vertices $u, v \in V \setminus U$ we have that the number of vertices $v^{(i)} \in U$ with $u \prec v^{(i)} \prec v$ is even.

Exercise 3.51 Show, using the evenness criterion from the previous exercise, that the following holds for the number $f_{n,m}$ of facets of a cyclic polytope $Z_n(m)$:

$$f_{n,m} = \begin{cases} \frac{m}{m - \frac{n}{2}} \binom{m - \frac{n}{2}}{m - n} & \text{if } n \text{ is even,} \\ 2\binom{m - \frac{n+1}{2}}{m - n} & \text{if } n \text{ is odd.} \end{cases} \tag{3.2}$$

Exercise 3.52 Compute the group of combinatorial automorphisms of each cyclic polytope.

In the remainder of this section we will discuss the relationship between the number of faces of varying dimensions of polytopes. These relations are essential for a deeper understanding of the combinatorics of polytopes (such as the proof of the exact statement of the Upper-bound Theorem).

The entries of the f-vector of a polytope are not independent of each other. This is easy to see for simple n-polytopes. Here, every vertex is incident with exactly n edges and conversely, every edge is incident with exactly two vertices. This implies $2f_1 = nf_0$. Since f_1 is an even number, this implies that every simple polytope of odd dimension has an even number of vertices. In the dual picture this means that each simplicial polytope of odd dimension has an even number of facets. Theorem 3.54 sharpens this statement. First we look at a famous result that holds for arbitrary polytopes.

Theorem 3.53 (Euler's formula) *The f-vector of a non-empty polytope P of dimension n satisfies the following equation*

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

Euler's formula implies that for two-dimensional polytopes the number of vertices and edges in a polygon is the same.

For three-dimensional polytopes we obtain the classical formula for the *Euler characteristic*:

$$f_0(P) - f_1(P) + f_2(P) = f_{-1}(P) + f_3(P) = 2. \tag{3.3}$$

Proof We prove Euler's formula by induction over the dimension n of the polytope.

For $n = 1$ each polytope has exactly two proper faces, i.e., its vertices, such that

$$\sum_{k=-1}^{1} (-1)^k f_k(P) = 1 - 2 + 1 = 0.$$

So let P be an n-polytope and let $m = f_0(P)$ be the number of vertices of P. After a suitable affine transformation we can assume without loss of generality that no two vertices have the same x_n-coordinates. Let $v^{(1)}, \ldots, v^{(m)}$ be the vertex set of P, ordered increasingly by their x_n-coordinate. Furthermore, let $H_1, \ldots, H_{2m-1}$ be horizontal (i.e., orthogonal to the x_n-axis) affine hyperplanes such that $v^{(i)} \in H_{2i-1}$, $1 \leq i \leq m$, and such that $v^{(i)}$ is the only vertex that is located between H_{2i-2} and H_{2i}. For a face F of P we define

$$\chi_j(F) = \begin{cases} 1 & \text{if } H_j \cap \operatorname{relint} F \neq \emptyset, \\ 0 & \text{otherwise} \end{cases}$$

for $1 \leq j \leq 2m - 1$.

Now we fix a face F and denote by $v^{(l)}$ the vertex with minimal x_n-coordinate. Similarly $v^{(u)}$ is the vertex with maximal x_n-coordinate. The horizontal hyperplanes that intersect the interior of F lie strictly between the hyperplanes H_{2l-1} and H_{2u-1}. If $\dim F \geq 1$ then we have $l \neq u$ and the number of hyperplanes with even index that intersect relint F exceeds the number of hyperplanes with odd index that intersect relint F by one. That is,

$$\sum_{j=2}^{2m-2} (-1)^j \chi_j(F) = 1.$$

Summing this equation over the set $\mathcal{F}_k(P)$ of k-faces of P yields

$$f_k(P) = \sum_{F \in \mathcal{F}_k(P)} \sum_{j=2}^{2m-2} (-1)^j \chi_j(F).$$

The alternating sum over all $k \geq 1$ yields

$$\sum_{k=1}^{n} (-1)^k f_k(P) = \sum_{j=2}^{2m-2} (-1)^j \sum_{k=1}^{n} (-1)^k \sum_{F \in \mathcal{F}_k(P)} \chi_j(F). \qquad (3.4)$$

For $2 \leq j \leq 2m-2$, $P_j := P \cap H_j$ has dimension $n-1$, so that by the induction hypothesis we have

$$\sum_{k=0}^{n-1} (-1)^k f_k(P_j) = 1. \qquad (3.5)$$

We distinguish between two cases:

j even: Each $(k-1)$-face of P_j is the intersection of a k-face of P with the hyperplane H_j such that

$$f_{k-1}(P_j) = \sum_{F \in \mathcal{F}_k(P)} \chi_j(F), \quad \text{for } 1 \leq k \leq n.$$

Substituting this into (3.5) yields

$$\sum_{k=1}^{n} (-1)^{k-1} \sum_{F \in \mathcal{F}_k(P)} \chi_j(F) = 1. \qquad (3.6)$$

j odd: Each $(k-1)$-face of P_j is the intersection of a k-face of P with H_j, with the exception of the vertex $v^{((j+1)/2)}$ which is contained in H_j. We therefore have

$$f_0(P_j) = 1 + \sum_{F \in \mathcal{F}_1(P)} \chi_j(F),$$

$$f_{k-1}(P_j) = \sum_{F \in \mathcal{F}_k(P)} \chi_j(F), \quad 2 \le k \le n.$$

In this case, substituting into (3.5) yields

$$\sum_{k=1}^{n} (-1)^{k-1} \sum_{F \in \mathcal{F}_k(P)} \chi_j(F) = 0. \tag{3.7}$$

Multiplying (3.6) and (3.7) by $(-1)^{j+1}$ and substituting into (3.4) yields

$$\sum_{k=-1}^{n} (-1)^k f_k(P) = -1 + m + \sum_{k=1}^{n} (-1)^k f_k(P)$$

$$= -1 + m + (m-1) \cdot (-1) + (m-2) \cdot 0 = 0. \qquad \square$$

From this we can deduce, with a clever summation, a previously mentioned, far-reaching generalization of the relation $2f_1 = nf_0$ for simple polytopes. For the proof we refer to Ziegler [99, §8.3].

Theorem 3.54 (Dehn–Sommerville equations) *The f-vector of a simple n-polytope P satisfies the linear equations*

$$\sum_{j=0}^{k} (-1)^j \binom{n-j}{n-k} f_j(P) = f_k(P), \quad \text{for } k \in \{0, \dots, n\}.$$

Through duality we obtain a corresponding statement for simplicial polytopes.

3.6 Inspection Using `polymake`

We want to study some concrete examples of polytopes, most of the characteristics of which can be obtained using the results discussed above. Here, and in the following, we use the software `polymake`, which is briefly introduced in Appendix D.1. The shell-based interface uses a dialect of Perl.

3.6.1 Cyclic Polytopes

We study cyclic 4-polytopes $Z_4(7)$ with 7 vertices. The `polymake` function `cyclic` generates cyclic polytopes which can then be further examined. Each line starting with "`polytope >`" contains one command.

```
polytope > $Z_4_7 = cyclic(4,7);

polytope > print $Z_4_7->VERTICES;
1 0 0 0 0
1 1 1 1 1
1 2 4 8 16
1 3 9 27 81
1 4 16 64 256
1 5 25 125 625
1 6 36 216 1296

polytope > print $Z_4_7->DIM;
4

polytope > print $Z_4_7->F_VECTOR;
7 21 28 14

polytope > print dense($Z_4_7->VERTICES_IN_FACETS);
1 1 0 0 1 1 0
0 1 1 0 1 1 0
0 0 1 1 1 1 0
0 1 1 1 1 0 0
1 1 0 1 1 0 0
1 1 1 1 0 0 0
1 1 1 0 0 0 1
1 0 1 1 0 0 1
1 0 0 1 1 0 1
0 0 0 1 1 1 1
1 0 0 0 1 1 1
0 0 1 1 0 1 1
0 1 1 0 0 1 1
1 1 0 0 0 1 1
```

VERTICES, DIM, F_VECTOR and VERTICES_IN_FACETS are examples of
properties of a `polymake` object. Each line of the property VERTICES contains
the oriented homogeneous coordinates of a vertex. The vertices are implicitly enu-
merated, starting with 0, but the order is not relevant here. The output of VER-
TICES_IN_FACETS refers to this enumeration: Every line of the output matrix
corresponds to a facet and every column to a vertex, where the order of the columns
corresponds to the order of the vertices in the VERTICES section. A 1 at position
(i, j) indicates that the i-th facet is incident with the j-th vertex. If the `dense()`
command is omitted in the `polymake`-command line, we get the entire list of ver-
tices for each facet.

   In the dense output of VERTICES_IN_FACETS we can immediately verify
Gale's evenness criterion from Exercise 3.50: In each line we have an even number
of 1s between two 0s. The matrix with coefficients in {0, 1} coded in the property
VERTICES_IN_FACETS is called the *incidence matrix* with respect to the given
order of vertices and facets.
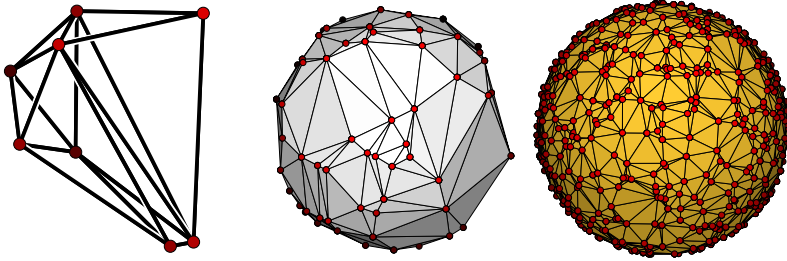
Fig. 3.10 The convex hulls of 8, 100 and 1000 random points on $\mathbb{S}^2$

The remaining two commands print the dimension and the f-vector of $Z_4(7)$. In particular, we see that this polytope has seven vertices, 21 edges, 28 ridges and 14 facets.

The homogeneous coordinates of the facets can be obtained using the command

```
polytope > print $Z_4_7->FACETS;
```

in the same order as for the property VERTICES_IN_FACETS. The output (which is suppressed here) looks similar to that in Section 5.4.

3.6.2 Random Polytopes

The function `rand_sphere` produces n-polytopes as convex hulls of random points selected from a uniform distribution on the unit sphere $\mathbb{S}^{n-1}$. By Exercise 3.20 these points are almost certainly in general position so that the convex hull is simplicial.

```
polytope > $R_3_8 = rand_sphere(3,8);
polytope > print $R_3_8->SIMPLICIAL;
1
polytope > print $R_3_8->F_VECTOR;
8 18 12
```

By the Dehn–Sommerville equations we know that the complete f-vector of a simplicial 3-polytope is determined by the number of vertices. We have $f_2 = 2f_0 - 4$ and $f_1 = f_0 + f_2 - 2 = 3f_0 - 6$. Figure 3.10 depicts some random polytopes.

3.6.3 Projective Transformations

We now want to show how to use `polymake` to projectively transform an unbounded but pointed polyhedron into a polytope. This is an example of the procedure described in Theorem 3.36. Let us begin by defining an unbounded polyhedron

as the Minkowski sum of a polytope (which is the convex hull of eight points) and
one infinite ray.

```
polytope > $P=new Polytope(POINTS=>
    [[0,0,0,1],
     [1,0,0,0],[1,3,0,0],[1,0,3,0],[1,3,3,0],
     [1,1,1,-1],[1,1,2,-1],[1,2,1,-1],[1,2,2,-1]]);
```

Again we use homogeneous coordinates, and the unique ray is represented by the
vector $(0, 0, 0, 1)^T$ listed first. It is easy to verify that $P satisfies our conditions.

```
polytope > print $P->BOUNDED, " ", $P->POINTED;
0 1
```

The first step is to set up an affine transformation which sends our polyhedron into
the positive orthant. To this end we list which facet is incident with which vertex.

```
polytope > print rows_numbered($P->FACETS_THRU_VERTICES);
0:1 2 5 6
1:0 1 2 3
2:0 1 6 8
3:2 3 5 7
4:5 6 7 8
5:0 3 4
6:3 4 7
7:0 4 8
8:4 7 8
```

In our input above the POINTS defining $P were, in fact, a non-redundant descrip-
tion. Hence the row numbers correspond to our input. For instance:

```
polytope > print $P->VERTICES->[5];
1 1 1 -1
```

We see that the facets numbered 0, 3 and 4 are incident with the vertex numbered 5.
Our polyhedron is full-dimensional and hence those three facets must have linearly
independent facet normal vectors. Therefore we can use vertex number 5 as our
point z in the construction described in Section 3.4. In this way we form the matrix
T as follows.

```
polytope > $T=new Vector([1,0,0,0])/$P
    ->FACETS->minor([0,3,4],All);
polytope > print $T;
1 0 0 0
0 0 1 1
0 1 0 1
1 0 0 1
```

Notice that the operator / concatenates matrices row-wise. Similarly, the matrix B
can be built from standard constructions using row and column concatenations, the
latter being expressed via | .

```
polytope > $B=ones_vector(4)/(zero_vector(3)|unit_matrix(3));
polytope > print $B;
1 1 1 1
0 1 0 0
0 0 1 0
0 0 0 1
```

To transform our polytope we have to take into account that `polymake` uses row vectors to represent the points, and thus transformations operate on the right. We can use the matrices constructed above if we transpose.

```
polytope > $Q=transform($P,transpose($B*$T));
polytope > print $Q->BOUNDED;
1
```

The vertices of the transformed polytope $Q is as follows.

```
polytope > print rows_numbered($Q->VERTICES);
0:1 1/3 1/3 1/3
1:1 0 0 1/2
2:1 0 3/5 1/5
3:1 3/5 0 1/5
4:1 3/8 3/8 1/8
5:1 0 0 0
6:1 1/2 0 0
7:1 0 1/2 0
8:1 1/3 1/3 0
```

Observe that the vertex numbered 5, which is the image of the point z under the transformation, is the origin. All the vertices are contained in the positive orthant.

3.7 Exercises

Exercise 3.55 Show that two polytopes are combinatorially equivalent if and only if there exists an ordering of their vertices and facets such that their corresponding incidence matrices are equal.

   A polytope is called *cubical* if all of its proper faces are combinatorially equivalent to cubes. For cubical polytopes there is a statement which is analogous to Lemma 3.47. This is an observation of Gil Kalai.

Exercise 3.56 Show that the following inequality holds for the f-vector of a cubical polytope

$$f_1 + 2f_2 + 2^2 f_3 + \cdots + 2^{n-2} f_{n-1} \leq \binom{f_0}{2}.$$

Exercise 3.57 Show that, given an arbitrary full-dimensional polyhedron $P \subseteq \mathbb{R}^n$ with outer description $P = \bigcap_{i=1}^{m} H_i^+$, there exists a family of indices $i_0, i_1, \ldots, i_n$

such that $Q = H_{i_0}^+ \cap \cdots \cap H_{i_n}^+$ is projectively equivalent to an n-simplex. If we additionally assume that P is a polytope, can we always choose the hyperplanes in such a way that Q is also a polytope?

Exercise 3.58 Let $\pi : \mathbb{R}^{n+1} \to \mathbb{R}^n$ be the linear projection to the first n coordinates. Show that the image of a polytope under π is again a polytope.

Exercise 3.59 Let P be an n-polytope. Show that there exists for every k-face G of P a family of facets $F_1, \ldots, F_{n-k}$ such that

$$G_1 \supsetneq G_2 \supsetneq \cdots \supsetneq G_{n-k} = G$$

holds for $G_i := F_1 \cap \cdots \cap F_i$.

Exercise 3.60 The Minkowski-sum $[p^{(1)}, q^{(1)}] + \cdots + [p^{(k)}, q^{(k)}]$ of a finite number of line segments with $p^{(i)}, q^{(i)} \in \mathbb{R}^n$ is a *zonotope*. Show that the zonotopes generated by k segments are exactly the images of the standard cube $[-1, 1]^k$ under affine maps.

Exercise 3.61 Which among the polyhedra in the following list are projectively equivalent?

(a) $\mathrm{conv}\{(0, 0, 0)^T\} + \mathrm{pos}\{(1, 0, 0)^T, (0, 1, 0)^T, (0, 0, 1)^T\}$
(b) $\mathrm{conv}\{(0, 0, 0)^T, (1, 0, 0)^T, (0, 1, 0)^T\} + \mathrm{pos}\{(1, 1, 1)^T\}$
(c) $\mathrm{conv}\{(1, 0, 0)^T, (0, 1, 0)^T, (0, 0, 1)^T, (a, b, c)^T\}$ for a, b, c arbitrary real numbers (i.e., this is an infinite set of polyhedra)
(d) $\mathrm{conv}\{(1, 0, 0)^T, (0, 1, 0)^T, (0, 0, 1)^T, (2, 1, 1)^T, (1, 2, 1)^T, (1, 1, 2)^T\}$

3.8 Remarks

The content of this chapter forms part of the standard material for polytopes and polyhedra, see the monographs of Boissonnat and Yvinec [15], Brøndsted [16], Grünbaum [56] and Ziegler [99]. The Upper-bound Theorem was proved by Mc-Mullen [78]. Further proofs can be found in the books of Mulmuley [80] and Ziegler [99]. As a general reference we recommend the *Handbook of Discrete and Computational Geometry* [49].

The term *polyhedron* is not always used in the same way as we have defined it in this book. In particular topologists often use this term to describe a simplicial or polyhedral complex embedded in $\mathbb{R}^n$; it may also describe a triangulated manifold.

The set of vertices and edges of a 3-dimensional polytope can be interpreted as the set of vertices and edges of a planar graph on a sphere. Hence (3.3) is a special case of Euler's formula for planar graphs, see [2, Chapter 11]. In fact, Euler's formula generalizes to cell complexes. We will see a glimpse of this in Section 13.1 at the very end of this book.

Concerning the graph $\Gamma(P)$ of an n-polytope P introduced in Exercise 3.23, Warren M. Hirsch conjectured in 1957 that any two vertices of P can be connected by a path in $\Gamma(P)$ of at most $m - n$ edges, where m is the number of facets. This famous conjecture, the *Hirsch conjecture*, was disproved by Santos in 2010 [90].