

5 Duality of Polyhedra

Duality in convex polyhedra is a very interesting notion not only in the theory of polyhedra but also in polyhedral computation. Duality implies that two basic representation conversions between V-representation and H-representation of a polyhedron are essentially the same thing. Yet, in order to convert one to the other is sometimes tricky because there are certain assumptions under which any specific conversion can work.

5.1 Face Lattice

Let P be a convex polytope in \mathbb{R}^d . Each face F of P is a convex polytope again by definition. The *dimension* $\dim(P)$ of P is the maximum number of affinely independent points in P minus one. The number of k -dimensional faces of P is denoted by $f_k(P)$. By Theorem 3.14, $f_k(P)$ is finite. A k -dimensional face (polytope) is called simply *k -face* (*k -polytope*). For a d -polytope P , the vector

$$f(P) := (f_{-1}, f_0, f_1, \dots, f_d) \quad (5.1)$$

is called the *f-vector* of P . Clearly $f_{-1} = f_d = 1$.

The 0-faces of a d -polytope are called the *vertices*, the 1-faces the *edges*, the $(d-1)$ -faces the *facets*, and the $(d-2)$ -faces the *ridges*.

We denote by $\mathcal{F}(P)$ the finite set of all faces of P ordered by set inclusion. This is called the *face lattice* of P . Recall that a lattice is a partially ordered set (poset in short) where the join (the least upper bound) and the meet (the greatest lower bound) of any two elements a and b exist in the set. The face lattice of a polytope is also known as the *combinatorial structure* of the polytope. In Figure 5.1, the face lattices of 1-, 2- and 3-cubes are depicted, whose f-vectors are $\{1, 2, 1\}$, $\{1, 4, 4, 1\}$ and $\{1, 8, 12, 6, 1\}$. One can easily show that all 1-polytopes are finite line segments and thus are combinatorially the same *diamond*.

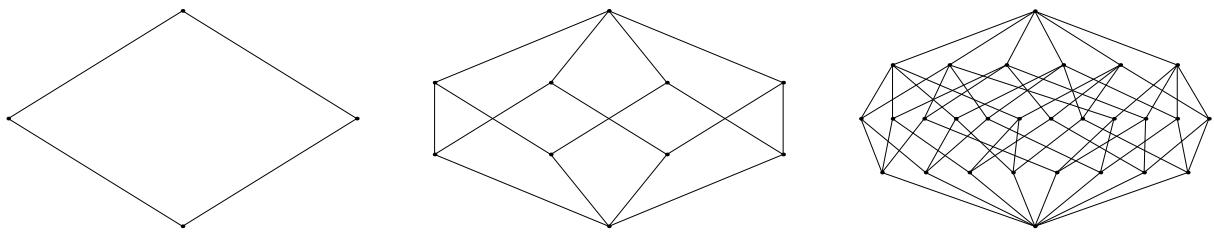


Figure 5.1: The Hasse diagram of the face lattices of 1-, 2- and 3-cubes

A lattice is called *polytopal* if it is isomorphic to the lattice of a polytope. Polytopal lattices are very special. The following proposition summarizes this.

Proposition 5.1 *Every polytopal lattice satisfies the following properties.*

(a) *It satisfies the Jordan-Dedekind chain property, i.e., all maximal chains between any two ordered elements $a < b$ have the same length.*

(a) *Fny two ordered elements $a < b$, the interval $[a, b]$, the set of elements between a and b , is again a polytopal lattice. In particular, this means that every interval of hight 2 is a diamond.*

We shall prove these properties in a later section, as we do not use them to prove the polytopal duality to be described below.

For a polytope P , a polytope P' is called a *dual* of P if $\mathcal{F}(P')$ is anti-isomorphic to $\mathcal{F}(P)$. It follows that a polytope P and a dual polytope P' have the same dimension, and their f-vectors are reversed, $f_i(P) = f_{d-i-1}(P')$ for all $i = -1, 0, \dots, d$. The following is the fundamental theorem of duality.

Theorem 5.2 *Every polytope admits a dual polytope.*

It is easy to see that a dual polytope is not unique. Any d -simplex is a dual of a d -simplex. A 3-cube is a dual of an octahedron but there are many geometrically different polytopes with isomorphic face lattices.

Yet, there is a simple construction of a dual polytope which is extremely useful, both theoretically and computationally. For a convex body C in \mathbb{R}^d containing the origin $\mathbf{0}$ in its interior, define its *polar* denoted by C^* as

$$C^* = \{y \in \mathbb{R}^d : x^T y \leq 1, \forall x \in C\}. \quad (5.2)$$

Theorem 5.3 *Let P be a polytope containing the origin $\mathbf{0}$ in its interior. Then its polar P^* is a dual polytope of P .*

5.2 Active Sets and Face Representations

As we learned from Theorem 3.9, every polyhedron has two representations, H-representation and V-representation. These two representations are closely linked to duality. Intuitively, by setting the transpose of an H-representation as a V-representation, we obtain a dual. This statement is in general incorrect and can be stated correctly with proper assumptions.

Let P be a polyhedron with an H-representation (A, b) and a V-representation (V, R) . Each row of (A, b) is denoted by (A_i, b_i) , representing the inequality $A_i x \leq b_i$. Each column of V and R is denoted by v_j and r_k , respectively, the j th vertex generator and the k th ray generator. We employ a little abuse of language here. An H-representation (A, b) is considered as the set of all its rows (A_i, b_i) , and similarly V (R) is considered as the set of all its columns v_j 's (r_k 's).

Let F be a non-empty face of P . An inequality (A_i, b_i) is called *active* at F if the inequality is satisfied with equality at all points in F . The set of all active inequalities is called the *active inequality* set at F .

Similarly, a vertex generator v_j is called *active* at F if $v_j \in F$. A ray generator r_k is called *active* at F if moving from any point on F along the direction r_k won't leave the polyhedron, i.e., $x + \theta r_k \in F$ for any $x \in F$ and $\theta \geq 0$. The pair (V', R') of sets of all active vertices and active rays are called the *active generator* sets at F . We extend the set inclusion for pairs of sets in the natural way, we define $(V'', R'') \subseteq (V', R')$ if and only if $V'' \subseteq V'$ and $R'' \subseteq R'$.

Active inequalities and generators are very important for representation of faces and face lattices.

Theorem 5.4 *Let P be a polyhedron with a V -representation (V, R) , and let F be a nonempty face of P . Then the active generator set pair (V', R') at F is a V -representation of F .*

Proof. Let (J, K) be the column index sets of (V', R') , namely, $V' = (v_j : j \in J)$ and $R' = (r_k : k \in K)$. Let

$$\overline{F} = \{x \in \mathbb{R}^d : x = V'\mu' + R'\lambda', \mu' \geq \mathbf{0}, \mathbf{1}^T \mu' = 1, \lambda' \geq \mathbf{0}\}.$$

We need to show $F = \overline{F}$. By definition of active generators, we have $F \supseteq \overline{F}$. For the converse inclusion, let $p \in F$ and suppose $p \notin \overline{F}$. Since $p \in P$,

$$p = V\mu + R\lambda \tag{5.3}$$

for some $\mu \geq \mathbf{0}$, $\mathbf{1}^T \mu = 1$ and some $\lambda \geq \mathbf{0}$. Because $p \notin \overline{F}$, we have $\mu_j > 0$ for $j \notin J$ or $\lambda_k > 0$ for $k \notin K$. Suppose there is $j \notin J$ such that $\mu_j > 0$. Then, $v_j \notin F$. Let (A, b) be an H -representation of P . Since $v_j \notin F$, there is an inequality (A_i, b_i) active at F such that $A_i v_j < b_i$. Since this is active at F , $A_i p = b_i$ and this implies that there is a ray or vertex generator in the RHS representation 5.3 of p which violates this active inequality. This is impossible. The second case is impossible by a similar argument. Thus, $p \in \overline{F}$. This completes the proof. \blacksquare

Theorem 5.5 *Let P be a polyhedron with an H -representation (A, b) and a V -representation (V, R) . Then*

- (a) *the face poset $\mathcal{F}(P) \setminus \{\emptyset\}$ is anti-isomorphic to the set of all active inequality sets ordered by set inclusion;*
- (b) *the face poset $\mathcal{F}(P) \setminus \{\emptyset\}$ is isomorphic to the set of all active generator sets ordered by set inclusion.*

Proof.

- (a) It is clear that the larger a face is, the smaller its active inequality set is. The main question is if the strictly larger a face is, the strictly smaller its active inequality set is. This follows directly from Theorem 3.14.
- (b) Using a similar argument to (a), (b) follows from Theorem 5.4. \blacksquare

5.3 Duality of Cones

Before proving the duality of polytopes, we show the duality of cones is a straightforward consequence of Theorem 5.5, a basic theorem on face lattice representations by active sets.

The notion of dual (polyhedral) cone is essentially the same as that of polytopes, the face lattice of a dual is the polar (upside-down) lattice. There is a small technical difference. Cones are different from polytopes in the sense that every cone has a unique minimal nonempty face (containing the origin), which plays exactly like the empty face of every polytope. For this reason, we define the *face lattice $\mathcal{F}(C)$ of a cone C* as the set of all **nonempty faces** of C ordered by set inclusion. Accordingly, we say that a cone C' is a *dual* cone of a cone C if $\mathcal{F}(C)$ and $\mathcal{F}(C')$ are anti-isomorphic.

Theorem 5.6 Every (polyhedral) cone admits a dual cone.

Our proof is by construction. For this, we define two cones. For a real $m \times d$ matrix A , we denote by $C_H(A)$ the cone with A as its H-representation:

$$C_H(A) = \{x : Ax \leq \mathbf{0}\}. \quad (5.4)$$

For a real $d \times s$ matrix R , we denote by $C_V(R)$ the cone with R as its V-representation:

$$C_V(R) = \{x : x = R\lambda, \lambda \geq \mathbf{0}\}. \quad (5.5)$$

Using this notation, Minkowski-Weyl Theorem, Theorem 3.10, says that a set C is of form $C = C_H(A)$ for some matrix A if and only if $C = C_V(A)$ for some matrix R .

The following is a stronger (constructive) version of the cone duality, Theorem 5.6.

Theorem 5.7 For any real $m \times d$ matrix A , the cone $C_H(A)$ and the cone $C_V(A^T)$ are dual to each other.

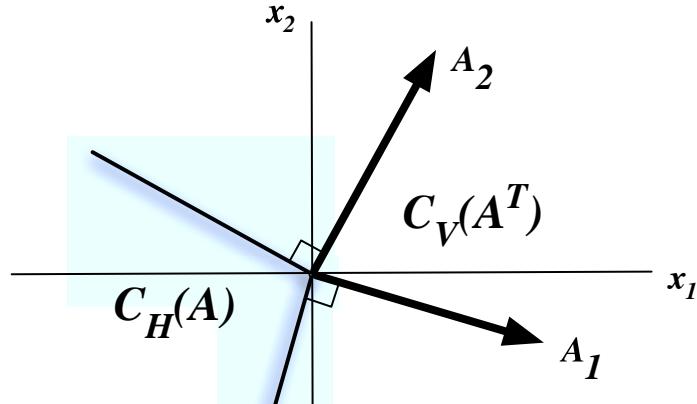


Figure 5.2: Cone Duality

Proof. Let A be a real $m \times d$ matrix. Let F be any nonempty face of $C_H(A)$, and let $I \subseteq [m]$ be the set of active inequality row indices at F , i.e., $F = \{x \in C_H(A) : A_Ix = \mathbf{0}\}$ and

$$\begin{aligned} \exists c \in \mathbb{R}^d \text{ such that } A_i c = \mathbf{0}, \forall i \in I, \text{ and} \\ A_j c < \mathbf{0}, \forall j \in [m] \setminus I. \end{aligned}$$

Or equivalently,

$$\begin{aligned} \exists c \in \mathbb{R}^d \text{ such that } c^T (A_i)^T = \mathbf{0}, \forall i \in I, \text{ and} \\ c^T (A_j)^T < \mathbf{0}, \forall j \in [m] \setminus I. \end{aligned}$$

Noting that the vectors $(A_i)^T$ ($i \in [m]$) are the generators of the cone $C_V(A^T)$, the relations above show exactly that $\{(A_i)^T : i \in I\}$ is the active generator set at the face of $C_V(A^T)$

determined by the valid inequality $c^T x \leq 0$ for the cone $C_V(A^T)$. The reverse direction is obvious.

This provides a one-to-one correspondence between the set of nonempty faces of $C_H(A)$ and the set of nonempty faces of $C_V(A^T)$, reversing the set inclusion. This completes the proof. \blacksquare

5.4 Duality of Polytopes

As we learned in the previous section, the duality of cones arises very naturally, and in fact, an H-representation of a cone immediately gives a V-representation of a dual cone, and vice versa: the cones $C_H(A)$ and $C_V(A^T)$ are dual to each other for any matrix A .

To derive a similar construction for polytopes, one can use the cone duality carefully. The main idea is to express a d -polytope P as the intersection of $(d+1)$ -cone C in such a way that P is embedded in C as the intersection of C with hyperplane $x_{d+1} = -1$. This is easy to do if P is a V-polytope. This gives some matrix R and the cone $C_V(R)$ in \mathbb{R}^{d+1} . We know how to construct a dual of $C_V(R)$: $C_H(R^T)$. The hard part is the rest: we have to make sure that it is “nicely” intersected by a hyperplane so that the intersection is a polytope and has the same face lattice as $C_H(R^T)$. If this is done, we have the construction of a dual of P .

Let P be a d -polytope with V-representation $V = [v_1, \dots, v_m]$. Let

$$\hat{V} := \begin{bmatrix} v_1 & v_2 & \cdots & v_m \\ -1 & -1 & \cdots & -1 \end{bmatrix}. \quad (5.6)$$

By Theorem 5.7, the following cones C and D defined below are dual to each other

$$C := C_V(\hat{V}) = \{x : x = \hat{V}\lambda, \lambda \geq \mathbf{0}\}, \quad (5.7)$$

$$D := C_H(\hat{V}^T) = \{x : \hat{V}^T x \leq \mathbf{0}\}. \quad (5.8)$$

Furthermore, by construction, the cone C represents P nicely.

Proposition 5.8 *The face lattices $\mathcal{F}(P)$ and $\mathcal{F}(C)$ are isomorphic.*

Proof. Consider the cut section P' of C with the hyperplane $h^{-1} := \{x \in \mathbb{R}^{d+1} : x_{d+1} = -1\}$. It follows that P and P' are affinely equivalent, and in particular, their face lattices are isomorphic. It is left to show that $\mathcal{F}(C)$ and $\mathcal{F}(P')$ are isomorphic. This follows from the fact that \hat{V} is not only a V-representation of C but also a V-representation of P' . \blacksquare

A proof of Theorem 5.2 is almost complete, because we know the face lattice of D is the target lattice we need to realize as the face lattice of a polytope. The only thing we have to show is that the cone D can be cut nicely by a hyperplane so that the cut section, say Q' , has the face lattice isomorphic to D . For this, consider the hyperplane $h^{+1} := \{x \in \mathbb{R}^{d+1} : x_{d+1} = +1\}$. define

$$Q' := D \cap h^{+1}. \quad (5.9)$$

Observe that

$$Q' = \left\{ \begin{bmatrix} x \\ x_{d+1} \end{bmatrix} \in \mathbb{R}^{d+1} : \begin{bmatrix} v_1^T & -1 \\ \vdots & \vdots \\ v_m^T & -1 \end{bmatrix} \begin{bmatrix} x \\ x_{d+1} \end{bmatrix} \leq \mathbf{0} \right\} \cap \{x \in \mathbb{R}^{d+1} : x_{d+1} = +1\} \quad (5.10)$$

$$= \left\{ \begin{bmatrix} x \\ 1 \end{bmatrix} \in \mathbb{R}^{d+1} : v_i^T x \leq 1, \forall i = 1, \dots, m \right\}. \quad (5.11)$$

Thus, the polyhedron Q' is affinely equivalent to the polyhedron

$$Q = \{x \in \mathbb{R}^d : v_i^T x \leq 1, \forall i = 1, \dots, m\} = \{x \in \mathbb{R}^d : V^T x \leq \mathbf{1}\}. \quad (5.12)$$

The polyhedron Q (and Q') may not have the face lattice isomorphic to D in general. Construct a small example to show this fact. The following lemma gives a right assumption for duality to work.

Theorem 5.9 *If P contains the origin in its interior, the polyhedron Q is a polytope dual to P .*

Proof. Assume that P contains the origin in its interior. The only thing left to be shown is that the face lattices of the cone D and the polyhedron Q' are isomorphic. For this, it is sufficient to show that Q' is bounded and a V-representation of Q' is in fact a V-representation of D . (Figure 5.3 shows that the assumption is in fact necessary.)

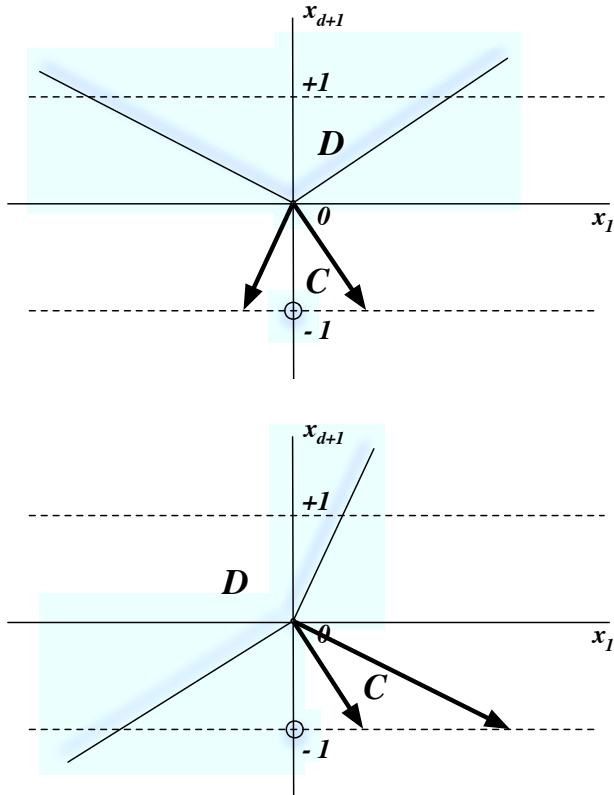


Figure 5.3: Polytope Duality: When it works and When it does not

Observe that the assumption is equivalent to

$$\text{rank } V = d, \text{ and} \quad (5.13)$$

$$\exists \lambda > \mathbf{0} \text{ such that } V\lambda = \mathbf{0}. \quad (5.14)$$

By (a variation of) Gordan's Theorem (Exercise 3.4), the statement (5.14) is equivalent to

$$\nexists x \text{ such that } V^T x \leq \mathbf{0}. \quad (5.14')$$

In order to show that Q (and Q') is bounded, suppose the contrary: there is an unbounded direction in Q , i.e., a nonzero vector z such that $V^T z \leq \mathbf{0}$. By (5.14'), this implies $V^T z = \mathbf{0}$ and $z \neq \mathbf{0}$, which is impossible by the assumption (5.14).

Now, we shall show a V-representation of Q' is a V-representation of D . For this, we take any nonzero vector $(x, x_{d+1})^T \in D$, and show that $x_{d+1} > 0$. This means that the normalized vector $(x/x_{d+1}, 1)^T$ is in Q' . Thus, any V-representation of Q' represents the cone D . To see that $x_{d+1} > 0$, observe that

$$\begin{aligned} \begin{bmatrix} x \\ x_{d+1} \end{bmatrix} \neq \mathbf{0} \text{ and } \hat{V}^T \begin{bmatrix} x \\ x_{d+1} \end{bmatrix} \leq \mathbf{0} \\ \implies \begin{bmatrix} x \\ x_{d+1} \end{bmatrix} \neq \mathbf{0} \text{ and } V^T x - \mathbf{1} x_{d+1} \leq \mathbf{0} \\ \implies x_{d+1} > 0. \end{aligned}$$

The last implication is valid because if $x_{d+1} \leq 0$, $V^T x \leq \mathbf{0}$ for $x \neq \mathbf{0}$ which is impossible by the assumptions (5.13) and (5.14). \blacksquare

5.5 Examples of Dual Pairs

In Section 3.7, we introduced a few examples of polytopes. Let's look at their duals.

First of all, one can easily see that a d -simplex is self-dual.

What is a dual of an d -cube? The simple way to see is to use the centrally symmetric cube $\text{Cube}(d)$ whose vertex set is $\{-1, 1\}^d$. Namely,

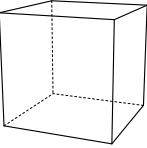
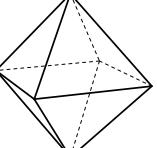
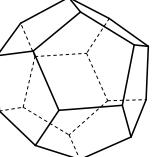
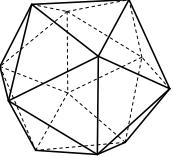
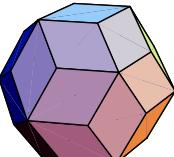
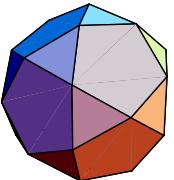
$$\text{Cube}(d) = \text{conv}\{-1, 1\}^d = \{x \in \mathbb{R}^d : \pm(e_i)^T x \leq 1, \forall i = 1, \dots, d\}. \quad (5.15)$$

The polar of $\text{Cube}(d)$ is known as the d -cross polytope:

$$\text{Cross}(d) = \{x : a^T x \leq 1, \forall a \in \{-1, 1\}^d\} = \text{conv}\{\pm e_i : i = 1, \dots, d\}. \quad (5.16)$$

Among all five regular polytopes in 3 dimension, the remaining duality is between a dodecahedron and an icosahedron.

An *icoso-dodecahedron* is a truncated dodecahedron, obtained from a dodecahedron truncated at each vertex to the midpoint of incident edges. The number of facets is clearly $30 = 12 + 20$, the sum of the numbers of facets of an icosahedron and a dodecahedron. Its dual is known as a *rhombic triacontahedron*, which is a very special zonotope arising as quasicrystal.

Type	Figure	# Vertices	# Facets	# i -Faces
Cube(d)		2^d	$2d$	$\binom{d}{i} 2^{d-i}$
Cross(d)		$2d$	2^d	$\binom{d}{i+1} 2^{i+1}$
Dodecahedron		20	12	
Icosahedron		12	20	
Rhombic Triacontahedron		32	30	
Icosa-Dodecahedron		30	32	

5.6 Simple and Simplicial Polyhedra

Both a 3-cube and a dodecahedron are *simple* polytopes and their duals are *simplicial* polytopes.

More generally, the *simple* d -polytopes are such that each vertex is contained in exactly d facets, while the *simplicial* d -polytopes are such that each facet contains exactly d vertices.

Proposition 5.10 *For a d -polytope P , the following statements are equivalent:*

- (a) *P is simple.*
- (b) *Each vertex v of P is incident to exactly d -edges.*
- (c) *For each vertex v of P , and for any k distinct edges incident to v , there exists a unique k -face containing the k edges.*
- (d) *For each vertex v of P , and for any 2 distinct edges incident to v , there exists a unique 2-face containing the 2 edges.*

Proposition 5.11 *For a d -polytope P , the following statements are equivalent:*

- (a) *P is simplicial.*
- (b) *Each facet f of P contains exactly d -ridges.*
- (c) *For each facet f of P , the intersection of any k distinct ridges contained in f is a $(d - k - 1)$ -face.*
- (d) *For each vertex v of P , the intersection of any 2 distinct ridges contained in f is a $(d - 3)$ -face.*

5.7 Graphs and Dual Graphs

Proposition 5.1 shows that every interval of height 2 is a diamond. This means one can define two types of graphs of a polytope. The *graph* of a polytope P is $G(P) = (V(P), E(P))$, where $V(P)$ is the set of vertices of P and $E(P)$ is the set of all edges each of which is represented as the pair of its two vertices. The *dual graph* of a polytope P is $G^D(P) = (F(P), R(P))$, where $F(P)$ is the set of facets of P and $R(P)$ is the set of all ridges each of which is represented as the pair of the two facets containing it. By the definition of duality, if Q is a dual of a polytope P , $G(P)$ is isomorphic to $G^D(Q)$.

6 Line Shellings and Euler's Relation

6.1 Line Shelling

Let $P = \{x \in \mathbb{R}^d : A_i x \leq 1, i = 1, 2, \dots, m\}$ be a polytope. P has such a representation iff it contains the origin in its interior.

A *shelling* of the boundary of P is a sequence F_1, F_2, \dots, F_m of its facets such that $(\cup_{i=1}^{k-1} F_i) \cap F_k$ is a topological $(d-2)$ -ball for each $2 \leq k \leq m-1$.

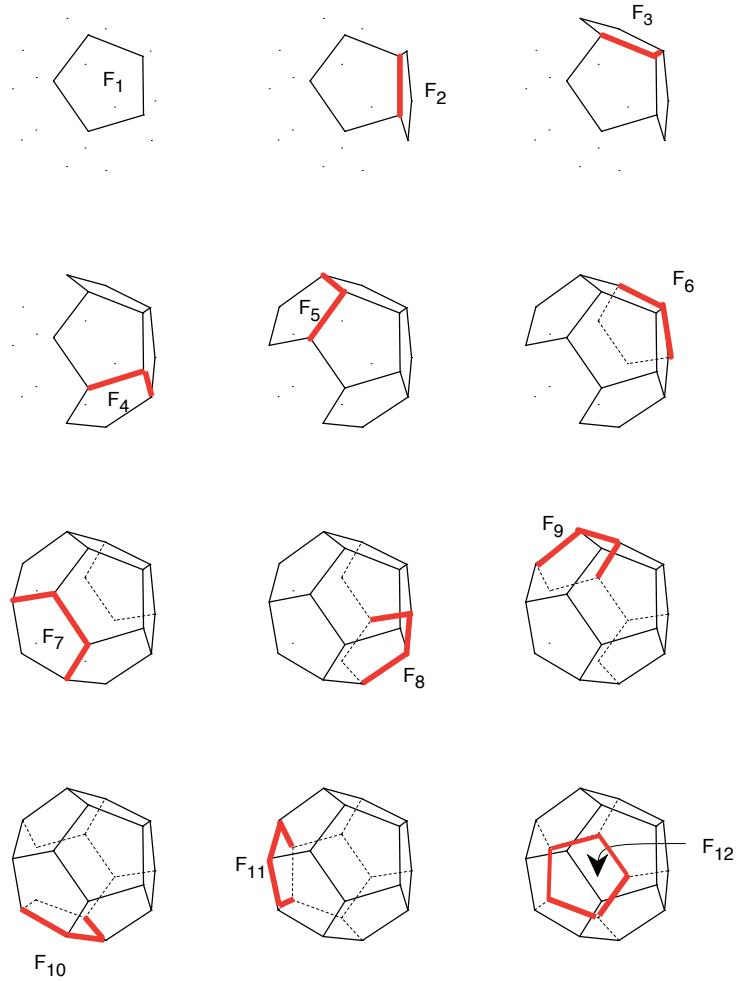


Figure 6.1: A Shelling of a Dodecahedron

The following is a fundamental theorem on polytopes which is extremely useful both theoretically and computationally.

Theorem 6.1 (Bruggesser-Mani [11] (1971)) *The boundary of every polytope admits a shelling.*

This theorem was used without proof by a Swiss mathematician Ludwig Schläfli (1901) to compute the Euler characteristic (to be discussed later) of convex polytopes. Seventy years later, an elegant proof was given. Here is the main idea. Bruggesser-Mani [11] proved a stronger theorem where any line in general position through an interior point of a polytope induces a particular shelling, known as a line shelling. Figure 6.2 illustrates this.

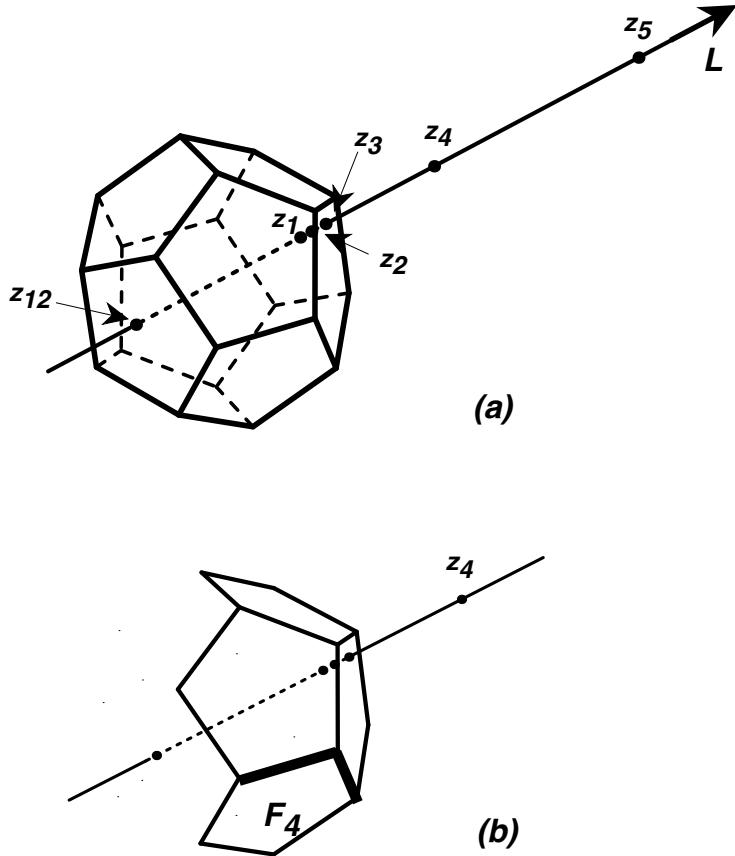


Figure 6.2: A Line Shelling

Imagine that a given polytope P is a planet earth and you are traveling along a generic oriented line L starting from some interior point. The first point z_1 to meet the boundary of P is a point on a facet. Let's call this facet F_1 . Then you meet another point z_2 on the boundary of a halfspace spanned by a facet. Let's call this facet F_2 . If you move a little forward from z_2 , you "see" only two facets F_1 and F_2 . This travel induces an ordering of facets as they become visible to you one by one. These facets are not all, and in the figure above, we have a sequence from F_1 up to F_6 . The rests are to be ordered in a similar manner but from the opposite side of infinity on L . More precisely, you travel from the other side of infinity and follow the line along the orientation. From a point far from P , you see all facets not yet ordered. Now moving toward P , some facet becomes invisible. Let's call this facet

F_7 . Then another facet becomes invisible that is F_8 . When we reach the last facet F_{12} , all facets of P are ordered. Such an ordering of facets can be shown to be a shelling of P .

For any point z in $\mathbb{R}^d \setminus P$, the union of all facets visible from z forms a very special subset of the boundary. Let's call it the *visible hemisphere* from z . Similarly, we define the *invisible hemisphere* from z . The proof uses the fact that both the visible and the invisible hemispheres are shellable.

Before giving a formal proof of Theorem 6.1, let us interpret the line shelling in a dual polytope.

Consider a polytope P in \mathbb{R}^d which contains the origin $\mathbf{0}$ in its interior. Thus, its H-representation can be of form

$$P = \{x \in \mathbb{R}^d : Ax \leq \mathbf{1}\}.$$

for some $m \times d$ matrix A . The polar dual of P is

$$P^* = \text{conv}\{A_i^T : i = 1, \dots, m\},$$

where A_i is the i th row of A .

For a generic $c \in \mathbb{R}^d$, sort the vertices A_i^T 's of the dual polytope by the linear function $c^T x$. Namely, we suppose

$$A_1 c > A_2 c > \dots > A_m c.$$

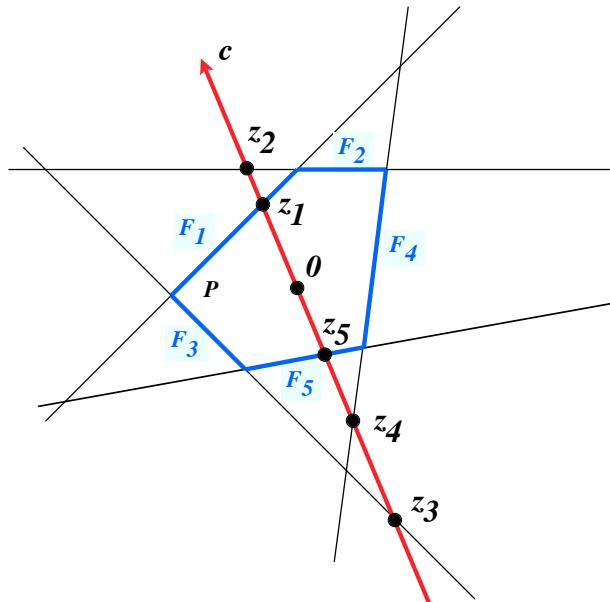
What is the meaning of this sorting for the original polytope P ?

Geometrically, the parameterized line $L(\lambda) = \{\lambda c \mid \lambda \in \mathbb{R}\}$ meets each hyperplane determined by $A_i x = 1$ at a point, say z_i . Let λ_i denotes the parameter value at the intersection. Thus,

$$z_i = \lambda_i c \quad \text{and} \quad A_i z_i = 1.$$

Consequently:

$$1/\lambda_1 > 1/\lambda_2 > \dots > 1/\lambda_k > 0 > 1/\lambda_{k+1} > \dots > 1/\lambda_m.$$



This ordering is exactly the ordering produced by the space travel. (For any positive $1/\lambda_i$, the smaller its value, the farther away the point z_i is from the origin. What about for negative case?)

6.2 Cell Complexes and Visible Hemispheres

A *cell complex* or simply a *complex* K in \mathbb{R}^d is a finite set of polytopes in \mathbb{R}^d such that

- (a) If $P \in K$ and F is a face of P , then $F \in K$.
- (b) If $P \in K$ and $Q \in K$, then $P \cap Q$ is a common face of P and Q .

The *dimension* $\dim K$ of a complex K is the largest dimension of its members. A complex of dimension d is called a *d-complex*. The *body* $|K|$ of a complex K is the union of all members of K . The *boundary complex* ∂K of a complex K is defined as the subcomplex of K consisting of all elements in K contained in the boundary of its body.

The *complex* $K(P)$ of a polytope P is the set of all faces of P . The boundary complex $\partial K(P)$ is simply the set of all proper faces of P . Both the complex and the boundary complex of a polytope are *pure*, i.e., the maximal members have the same dimension.

A pure complex K is called *B-shellable* if the maximal members can be arranged in a sequence F_1, F_2, \dots, F_m in such a way that the subcomplex induced by $(\cup_{i=1}^{k-1} F_i) \cap F_k$ is B-shellable for each $2 \leq k \leq m$. By definition, the complex of a 0-polytope is B-shellable, and those are the only B-shellable 0-complexes.

A pure complex K is called *S-shellable* if the maximal members can be arranged in a sequence F_1, F_2, \dots, F_m in such a way that the subcomplex induced by $(\cup_{i=1}^{k-1} F_i) \cap F_k$ is B-shellable (S-shellable, respectively) for each $2 \leq k \leq m-1$ ($k = m$). By definition, the boundary complex of a 1-polytope is S-shellable, and those are the only S-shellable 0-complexes.

These notions B-shellability and S-shellability are motivated by topological notions of balls and spheres. However, it should be observed that a B-shellable (S-shellable) complex is not necessarily a ball (a sphere). For example, the complex consisting of three 1-polytopes having a single vertex in common is B-shellable but not homeomorphic to a ball. We can add extra conditions to B-shellability (S-shellability) to enforce the resulting complexes to be topologically a ball (a sphere). The following is a combinatorial analogue of Theorem 6.1.

Theorem 6.2 *The boundary complex $\partial K(P)$ of a polytope is S-shellable.*

Before presenting a proof, we will give a nice application of this theorem. We define the *Euler characteristic* of a complex K as

$$\chi(K) = \sum_{i=0}^{\dim K} (-1)^i f_i(K), \quad (6.1)$$

where $f_i(K)$ is the number of i -dimensional members of K . It is easy to see that, for any two subcomplexes A and B of a complex, we have

$$\chi(A \cup B) + \chi(A \cap B) = \chi(A) + \chi(B). \quad (6.2)$$

Theorem 6.3 (Euler's Relation) *The following statements hold.*

- (a) *If K is B-shellable, $\chi(K) = 1$.*
- (b) *If K is S-shellable, $\chi(K) = 1 + (-1)^{\dim K}$.*

Proof. Both statements are obvious if $\dim K = 0$. By induction, we assume that both statements are true if $\dim K < d$ (≥ 1). First consider a B-shellable d -complex K . Since K is B-shellable, its d -polytopes can be ordered F_1, F_2, \dots, F_m in such a way that the subcomplex induced by $(\bigcup_{i=1}^{k-1} F_i) \cap F_k$ is B-shellable for each $2 \leq k \leq m$. When $m = 1$, clearly we have $\chi(K) = \chi(\partial K) + (-1)^d$. Since ∂K is S-shellable by Theorem 6.2 and has dimension $d - 1$, the induction hypothesis implies

$$\chi(K) = \chi(\partial K) + 1 = 1 + (-1)^{d-1} + (-1)^d = 1.$$

Now we use the second induction on m . We assume that (a) is valid if $f_d(K) < m$, and then consider the case $f_d(K) = m$. Since the subcomplex A induced by $(\bigcup_{i=1}^{m-1} F_i)$ is B-shellable by the second induction, it satisfies (a). We denote by B the subcomplex induced by F_m . By using the fact that the subcomplex C induced by $(\bigcup_{i=1}^{m-1} F_i) \cap F_m$ is a B-shellable $(d - 1)$ -complex, by the first induction and (6.2), we have

$$\chi(K) = \chi(A) + \chi(B) - \chi(C) = 1 + 1 - 1 = 1.$$

The remaining proof of (b) is straightforward as we already established (a). Let K be a S-shellable d -complex. Then, its d -polytopes can be ordered F_1, F_2, \dots, F_m such that the subcomplex induced by $(\bigcup_{i=1}^{k-1} F_i) \cap F_k$ is B-shellable (S-shellable, respectively) for each $2 \leq k \leq m - 1$ ($k = m$). Note that the subcomplex A induced by $(\bigcup_{i=1}^{m-1} F_i)$ is B-shellable and satisfies (a). The subcomplex induced by F_m is also B-shellable and satisfies (a). The subcomplex C induced by $(\bigcup_{i=1}^{m-1} F_i) \cap F_m$ is a S-shellable $(d - 1)$ -complex, by the first induction and (6.2), we have

$$\chi(K) = \chi(A) + \chi(B) - \chi(C) = 1 + 1 - (1 - (-1)^{d-1}) = 1 + (-1)^d.$$

This completes the proof. ■

Of special interest are topological properties of visible and invisible hemispheres of a polytope P . Please recall that for any point z in general position in $\mathbb{R}^d \setminus P$, the union of all facets of P is the *visible hemisphere* from z , denoted by $\text{vi}(P, z)$. Similarly, we define the *invisible hemisphere* from z , denoted by $\text{iv}(P, z)$.

Theorem 6.4 *Let P be a d -polytope in \mathbb{R}^d for $d \geq 1$ and let z be any point in general position in $\mathbb{R}^d \setminus P$. Then the two subcomplexes of $K(P)$ induced by the visible hemisphere $\text{vi}(P, z)$ and the invisible hemisphere $\text{iv}(P, z)$ from z are B-shellable.*

Proof. We use induction on d . By inductive hypothesis, we assume that (*) the visible hemisphere $\text{vi}(P, z)$ and invisible hemisphere $\text{iv}(P, z)$ from z induce B-shellable subcomplexes when $d < k$, with $k \geq 2$. The statement (*) is obvious for $d = 1$.

Now we try to show that $(*)$ is true for $d = k$. We take an oriented line L through z in general position which intersects the interior of P , and let z_1, z_2, \dots, z_m be the distinct intersections of L and the hyperplanes spanned by the facets F_1, \dots, F_m . Without loss of generality, the ordering is the one obtained by the space travel on L .

We first show that the visible hemisphere $\text{vi}(P, z)$ induces a B-shellable subcomplex. The point z is between z_i and z_{i+1} for some $1 \leq i < m$. If $i = 1$, $\text{vi}(P, z) = F_1$ and thus obviously $\text{vi}(P, z)$ induces a B-shellable subcomplex. We use induction again, on i , and assume by induction hypothesis that $\text{vi}(P, z)$ induces a B-shellable subcomplex for $i < h$ for some $h \geq 2$. We consider the case $i = h$. Note that $\text{vi}(P, z) = \text{vi}(P, z_i) \cup F_i$, where $\text{vi}(P, z_i)$ induces a B-shellable subcomplex by the inductive hypothesis. Now, we claim that $\text{vi}(P, z_i) \cap F_i$ induces a B-shellable $(d-2)$ -subcomplex. This is true because this set is in fact the visible hemisphere $\text{vi}(F_i, z_i)$ from z_i in the $(d-1)$ -dimensional space spanned by F_i . Since F_i is a $(d-1)$ -ball and $\text{vi}(P, z_i) \cap F_i$ induces a B-shellable subcomplex, $\text{vi}(P, z)$ induces a B-shellable subcomplex. Essentially the same argument shows that the invisible hemisphere $\text{iv}(P, z)$ induces a B-shellable subcomplex.

This completes the double induction proof. ■

Proof. (of Theorem 6.2) By definition, the boundary complex of any 1-polytope is S-shellable. We assume by induction that $\partial K(P)$ of any polytope of dimension $k-1$ or less is shellable. Consider any k -polytope P . Let F be a facet of P , and let z be a point from which F is the only visible facet of P . This means that $\text{iv}(P, z)$ is a subcomplex of $\partial K(P)$ induced by all facets of P except F . By Theorem 6.4, $\text{iv}(P, z)$ is B-shellable. We claim that any shelling ordering of $\text{iv}(P, z)$ with F appended at the end is a shelling of $\text{iv}(P, z)$. For this, we only need to show that $\partial K(F)$ is S-shellable. Since F has dimension $k-1$, this follows from the inductive hypothesis. This completes the proof. ■

6.3 Many Different Line Shellings

The proof of shellability of polytope boundaries using the notion of line shelling provides many different ways to shell a polytope boundary. The choice of a line is restricted only by the two conditions (1) it has to intersect with the interior of the polytope, (2) it must intersect the hyperplanes spanned by the facets at distinct points.

Proposition 6.5 *The boundary of every polytope admits a shelling F_1, F_2, \dots, F_m with any one of the following prescribed conditions:*

- (a) *both F_1 and F_m can be prefixed arbitrarily.*
- (b) *all facets incident to a given vertex can be ordered earlier than any other facets.*
- (c) *all facets incident to a given vertex can be ordered later than any other facets.*

7 McMullen's Upper Bound Theorem

7.1 Cyclic Polytops and the Upper Bound Theorem

The *moment curve* in \mathbb{R}^d is the image of the real space \mathbb{R} by the function $m(t)$ defined by

$$m(t) := (t, t^2, t^3, \dots, t^d)^T. \quad (7.1)$$

The function $m(\cdot)$ is thus a parametric representation of the moment curve.

A *cyclic polytope* is the convex hull of n ($> d$) distinct points on the moment curve, that is, $\text{conv}\{m(t_1), m(t_2), \dots, m(t_n)\}$ for some $t_1 < t_2 < \dots < t_n$. The following is a basic property of the moment curve.

Proposition 7.1 *Any $(d+1)$ distinct points on the moment curve $m(t)$ are affinely independent.*

Proof. Suppose $m(t_1), m(t_2), \dots, m(t_{d+1})$ are affinely dependent for some $t_1 < t_2 < \dots < t_{d+1}$. Then they must lie in some hyperplane, and thus there is a linear equation

$$a_0 + a_1 x_1 + a_2 x_2 + \dots + a_d x_d = 0$$

is satisfied by all $m(t_i)$'s. It follows that the polynomial equation

$$a_0 + a_1 t^1 + a_2 t^2 + \dots + a_d t^d = 0$$

is satisfied by $(d+1)$ distinct values of t , which contradicts to the fundamental theorem of algebra. \blacksquare

Proposition 7.1 implies that the cyclic polytope $c(d, n)$ is a simplicial polytope and its dual is a simple polytope.

As we will see that for any fixed d and n , its combinatorial structure is unique. Thus, we will denote anyone of them by $c(d, n)$, and their duals by $c^*(d, n)$.

McMullen's upper bound theorem is one of the most important theorems in the theory of convex polytopes.

Theorem 7.2 (McMullen's Upper Bound Theorem [38] (1970)) *For any fixed d and n , the maximum number of j -faces of a d -polytope with n vertices is attained by the cyclic polytope $c(d, n)$ for all $j = 0, 1, \dots, d-1$. Equivalently, for any fixed d and n , the maximum number of j -faces of a d -polytope with n facets is attained by the dual cyclic polytope $c^*(d, n)$ for all $j = 0, 1, \dots, d-1$.*

There is an explicit formula for $f_j(c(d, n))$ for $j = 0, 1, \dots, d-1$. The following gives essentially a half of these formulas.

Lemma 7.3 *For any $d \geq 0$ and $n \geq d+1$,*

(a) $f_{j-1}(c(d, n)) = \binom{n}{j}$, for $0 \leq j \leq \lfloor \frac{d}{2} \rfloor$.

(b) $f_k(c^*(d, n)) = \binom{n}{d-k}$, for $\lceil \frac{d}{2} \rceil \leq k \leq d$.

Proof. By duality, two statements (a) and (b) are equivalent. Let's prove (a).

Consider the cyclic polytope $P = \text{conv}\{m(t_1), m(t_2), \dots, m(t_n)\}$ with $t_1 < t_2 < \dots < t_n$. Let $0 \leq j \leq \lfloor \frac{d}{2} \rfloor$. Take the first j points $m(t_1), m(t_2), \dots, m(t_j)$, and consider the hyperplane h determined by

$$a_0 + a_1 x_1 + \dots + a_d x_d = 0, \quad (7.2)$$

where the coefficients a_i 's coincide with those in the polynomial

$$p(t) := a_0 + a_1 t + \dots + a_d t^d \equiv \prod_{i=1}^j (t - t_i)^2. \quad (7.3)$$

Note that the assumption $j \leq \lfloor \frac{d}{2} \rfloor$ implies that the polynomial $p(t)$ has degree at most d . Observe that h contains all the points $x(t_i)$ for $i = 1, \dots, j$. Furthermore, the remaining points $x(t_i)$ for $i = j+1, \dots, n$ are strictly on the positive side of the hyperplane. This means that $\text{conv}\{x(t_1), \dots, x(t_j)\}$ is a face of P . Since the above discussion works exactly the same way if we take any j points, every j points from $\{m(t_1), m(t_2), \dots, m(t_n)\}$ determine a $(j-1)$ -face. \blacksquare

Lemma 7.3 implies an interesting property of the cyclic polytope. Namely, if $d \geq 4$, then every pair of vertices forms an edge. This means that the graph of $c(d, n)$ is a complete graph for $d \geq 4$. This is not very intuitive because this phenomenon does not occur in the 3-dimensional space.

The proof of Lemma 7.3 gives some ideas on how to determine the facets of $c(d, n)$. Which d -tuples of points from $\{m(t_1), m(t_2), \dots, m(t_n)\}$ span a facet? Since any d -tuple $\{m(t_{j_1}), m(t_{j_2}), \dots, m(t_{j_d})\}$ is affinely independent and thus spans a hyperplane. Whether or not it defines a facet is thus equivalent to whether or not all the remaining points are on one side of the hyperplane. This turns out to be quite easy to check through a combinatorial condition, known as *Gale's evenness condition*.

Exercise 7.1 Find a necessary and sufficient condition for a set of d points $m(t_{j_1}), m(t_{j_2}), \dots, m(t_{j_d})$ to determine a facet of the cyclic polytope.

Lemma 7.3 gives essentially a half of the f -vector of the cyclic polytope. Yet, by using the fact that it is simplicial, the remaining information on the f -vector will be shown to be determined uniquely.

7.2 Simple Polytopes and h -vectors

To establish the Upper Bound Theorem, Theorem 7.2, we shall prove the dual statement:

For any fixed d and n , the maximum number of j -faces of a d -polytope with n facets is attained by the dual cyclic polytope $c^*(d, n)$ for all $j = 0, 1, \dots, d-1$.

We have two basic steps. **The first step** is to show that it is sufficient to consider only simple polytopes as maximizers of the number of j -faces, for a fixed number of facets. More precisely, we have:

Theorem 7.4 *For any d -polytope P with n facets in \mathbb{R}^d , there exists a simple d -polytope P' with n facets such that $f_j(P) \leq f_j(P')$ for all $j = 0, 1, \dots, d-1$.*

Proof. We only have to argue that a small perturbation of each inequality defining P does not decrease the number of j -faces. We leave the proof to the reader. Use Theorem 3.14 and analyze how a face changes as an inequality gets perturbed slightly toward enlarging the polytope. \blacksquare

The second step is to show that among all simple d -polytopes with n facets, the dual cyclic polytope $c^*(d, n)$ maximizes the number of j -faces for all $j = 0, 1, \dots, d-1$.

For the rest of this section,

(*) we only consider simple d -polytopes with n facets.

We denote by $\square(d, n)$ the set of all simple d -polytopes in \mathbb{R}^d with n facets.

For any $P \in \square(d, n)$, consider a linear program $c^T x$ subject to $x \in P$. Assume that c is generic and in particular, the LP orientation $\vec{G}(P)$ of the graph $G(P)$ is well-defined, namely, $\vec{G}(P)$ is a directed graph with the unique sink (maximizer) vertex and the unique source (minimizer) vertex.

Now, we denote by $h_k(\vec{G}(P))$ the number of vertices of indegree k , for each $k = 0, 1, \dots, d$. Clearly, $h_0(\vec{G}(P)) = h_d(\vec{G}(P)) = 1$. We shall eventually write $h_k(P)$ instead of $h_k(\vec{G}(P))$, as we will see below that this number does not depend on c at all.

Lemma 7.5 *For any polytope $P \in \square(d, n)$ and any generic $c \in \mathbb{R}^d$, the value $h_k(\vec{G}(P))$ depends only on P , and in particular, it does not depend on c . Thus, it can be denoted as $h_k(P)$.*

Proof. Let P be a polytope $P \in \square(d, n)$ and take generic $c \in \mathbb{R}^d$. We denote by (F, v) a pair of a face F of P and a vertex v on F which is the unique sink on F . It is clear that the number of such pairs (F, v) is the number of k -faces, $f_k(P)$.

Now, fix a vertex v of P and fix k . The number of such pairs (F, v) can be counted by using Proposition 5.10. Namely, there are exactly $\binom{r}{k}$ k -faces incident to v whose sink is v , where r is the indegree of v in $\vec{G}(P)$. Now ranging v over all vertices, we have

$$\sum_{r=0}^d h_r(\vec{G}(P)) \binom{r}{k} = f_k(P), \quad \text{for } k = 0, 1, \dots, d. \quad (7.4)$$

Now the system of linear equations can be written using a matrix and vectors as

$$\begin{bmatrix} \binom{0}{0} & \binom{1}{0} & & \cdots & \binom{d}{0} \\ 0 & \binom{0}{1} & \binom{2}{1} & & \cdots & \binom{d}{1} \\ 0 & 0 & \ddots & \cdots & & \vdots \\ 0 & 0 & 0 & \binom{k}{k} & \cdots & \binom{d}{k} \\ 0 & 0 & 0 & 0 & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & \binom{d}{d} \end{bmatrix} \begin{bmatrix} h_0 \\ h_1 \\ \vdots \\ h_k \\ \vdots \\ h_d \end{bmatrix} = \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_k \\ \vdots \\ f_d \end{bmatrix}. \quad (7.5)$$

The matrix on the LHS is obviously nonsingular, because it is upper triangular and the diagonal entries are all 1's. This means that h_j 's are determined uniquely by f_j 's, and thus h_j 's are independent of c . This completes the proof. \blacksquare

There are different ways to solve the equation (7.5) in terms of the h -vector $h(P) := (h_0, h_1, \dots, h_d)$. The resulting formula for h in terms of f is given by

$$h_i(P) = \sum_{k=0}^d (-1)^{k-i} \binom{k}{i} f_k(P), \quad \text{for } i = 0, 1, \dots, d. \quad (7.6)$$

This together with Lemma 7.3 provides us with simple explicit formulas for a half of the h -vector of the dual cyclic polytope.

Lemma 7.6

$$h_i(c^*(d, n)) = \binom{n-i-1}{d-i}, \quad \text{for } i = \lceil d/2 \rceil, \dots, d. \quad (7.7)$$

Proof. Substitute $f_k(P)$ in (7.6) with the explicit formula for $f_k(c^*(d, n))$ in Lemma 7.3 (b). Exercise. \blacksquare

The remaining part of the h vector comes for free, as we observe that the h -vector is symmetric, namely, by the definition of h_i ,

$$h_i(P) = h_{d-i}(P), \quad \text{for } i = 0, 1, \dots, \lfloor d/2 \rfloor, \quad (7.8)$$

where the RHS counts the LHS using the h -vector with the reversed orientation by the vector $-c$. These equations, expressed in terms of f -vector via (7.6), are known as the *Dehn-Sommerville Relations*.

Theorem 7.7 (The Dehn-Sommerville Relations) *Every simple d -polytope P satisfies the following equations.*

$$\sum_{k=i}^d (-1)^k \binom{k}{i} f_k(P) = \sum_{k=d-i}^d (-1)^{k-d} \binom{k}{d-i} f_k(P), \quad \text{for } i = 0, 1, \dots, \lfloor d/2 \rfloor. \quad (7.9)$$

More explicitly, the first two equations are

$$\sum_{k=0}^d (-1)^k f_k(P) = f_d(P) = 1, \quad (\text{i.e. Euler's Relation}), \quad (7.10)$$

$$-f_1(P) + 2f_2(P) - 3f_3(P) + \dots + (-1)^d d f_d(P) = -f_{d-1}(P) + d f_d(P). \quad (7.11)$$

The equation (7.5) shows that each f_j is a nonnegative combination of h_j 's. Therefore, the following is a strengthening of the Upper Bound Theorem, saying that the h -vector is component-wise maximized by the dual cyclic polytope.

Theorem 7.8 (A Strengthened Upper Bound Theorem) *For any simple d -polytope P with n facets the following inequalities hold.*

$$h_i(P) \leq h_i(c^*(d, n)), \quad \text{for } i = 0, 1, \dots, d. \quad (7.12)$$

Proof. Let P be a simple d -polytope with n facets. The claimed inequalities are trivial for $i = 0$ and $i = d$. By the symmetry of the h -vector, we only need to show

$$h_i(P) \leq h_i(c^*(d, n)) \equiv \binom{n-i-1}{d-i}, \quad \text{for } i = \lceil d/2 \rceil, \dots, d.$$

We use induction on i but with decreasing values. Suppose the theorem is valid for $i = k+1$ ($k < d$), and consider the case $i = k$.

We claim two inequalities for h -vectors. First we observe that for any facet F of P and for any i ,

$$h_i(F) \leq h_{i+1}(P). \quad (7.13)$$

This is valid because we can select a generic c such that all the vertices in F take the object value higher than any other vertices of P . Note that the values $h_i(F)$ and $h_{i+1}(P)$ are invariant over choices of c . Secondly, we have

$$\sum_F h_i(F) = (i+1)h_{i+1}(P) + (d-i)h_i(P). \quad (7.14)$$

The summation in LHS is over all faces F of P . This equation can be verified once we observe that every vertex of a face with indegree i in the face has indegree i or $i+1$ in P . If it has indegree i in P , there are exactly $(d-i)$ facets containing it that preserve the same indegree. If it has indegree $i+1$ in P , there are exactly $(i+1)$ facets containing it that decrease its indegree by one.

Now we look at the inductive step for $i = k$. By the two inequalities (7.14) and (7.13), we have

$$(k+1)h_{k+1}(P) + (d-k)h_k(P) = \sum_F h_k(F) \leq nh_{k+1}(P). \quad (7.15)$$

This implies

$$(d-k)h_k(P) \leq (n-k-1)h_{k+1}(P), \quad \text{or equivalently,} \quad (7.16)$$

$$h_k(P) \leq \frac{n-k-1}{d-k}h_{k+1}(P). \quad (7.17)$$

Now we use the inductive hypothesis for $i = k+1$ to get

$$h_k(P) \leq \frac{n-k-1}{d-k}h_{k+1}(P) \leq \frac{n-k-1}{d-k} \binom{n-k-2}{d-k-1} = \binom{n-k-1}{d-k}. \quad (7.18)$$

This completes the proof. ■

While the h -vector of the cyclic polytope is extremely simple, its f -vector is rather complicated. The formula can be written explicitly using (7.5), (7.6) and (7.8). We here present a formula for $f_0(c^*(d, n))$ which is quite simple.

Theorem 7.9 *The maximum number of vertices a d -polytope with n facets can have is realized by the dual cyclic polytope and is*

$$f_0(c^*(d, n)) = \binom{n - \lceil d/2 \rceil}{n - d} + \binom{n - \lfloor d/2 \rfloor - 1}{n - d}. \quad (7.19)$$

By duality, this number coincides with $f_{d-1}(c(d, n))$.

Proof. Left to the reader. Hint: use the identity:

$$\binom{n}{0} + \binom{n+1}{1} + \cdots + \binom{n+s}{s} = \binom{n+s+1}{s}.$$

■

8 Basic Computations with Polyhedra

Consider a system of m linear inequalities in d variables

$$Ax \leq b. \quad (8.1)$$

An inequality $A_i x \leq b_i$ is called *redundant* in (8.1) if the set of solutions to (8.1) stays unchanged when the inequality is removed from the system. An equivalent condition is that there is no x satisfying $A_i x > b_i$ and $A_j x \leq b_j$ for all $j \neq i$.

In this section, we study basic problems in polyhedral computation such as the following two problems:

Problem 8.1 [Single H-Redundancy]

Input: A rational matrix $A \in \mathbb{Q}^{m \times d}$, a rational vector $b \in \mathbb{Q}^m$ and an index $k \in [m] := \{1, \dots, m\}$

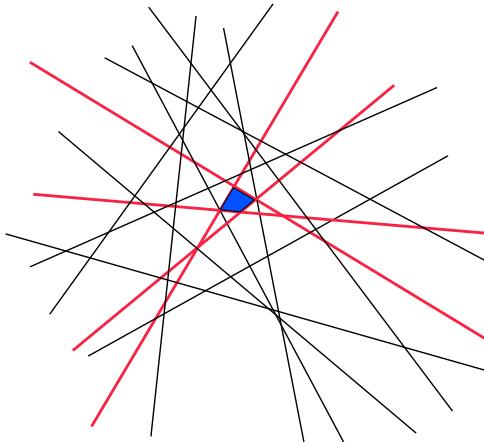
Output: **Yes** if $A_k x \leq b_k$ is redundant in $Ax \leq b$, **No** otherwise.

Problem 8.2 [H-Redundancy Removal]

Input: A rational matrix $A \in \mathbb{Q}^{m \times d}$, a rational vector $b \in \mathbb{Q}^m$

Output: An equivalent subsystem of $Ax \leq b$ which is free of redundancies.

The second problem can be solved by solving the first problem for each inequalities, but interestingly, one can do better than that by dynamically selecting the ordering of inequalities to be processed.



The figure above illustrates the H-redundancy problem. The blue region is the feasible region $P = \{x : Ax \leq b\}$. The output of the computation is the set of inequalities indicated in red that are essential in the H-representation. Often, the size of output is much smaller than the size of input.

Naturally one can pose the same problems for V-polyhedra. It turns out that those problems can be reduced to the H-redundancy problems. We will see that the H-redundancy problems can be reduced further to the H-redundancy problems for the special case of H-cones. These transformations are discussed in Section 8.4.

Here is a closely related problem that should be solved before the H-Redundancy Removal is solved.

Problem 8.3 [H-Dimension]

Input: A rational matrix $A \in \mathbb{Q}^{m \times d}$, a rational vector $b \in \mathbb{Q}^m$

Output: The dimension of the polyhedron $P = \{x : Ax \leq b\}$.

A typical algorithm for this computes not only the dimension of P , but also a relative interior point of P , see Section 8.3. One can then embed the polytope in a lower-dimensional space so that P becomes full-dimensional.

These problems are much easier than other problems in polyhedral computation such as the representation conversion between V- and H-representations and computing the volume of a polytope. In fact, the problems discussed in this section are all polynomially solvable in the size of input.

The main goal of this section is to present many algorithms which are not only polynomial-time but also best possible in terms of the number of LP's that must be solved, or of the size of LP's that must be solved when the number of LP's to be solved is fixed.

For this purpose, we use the notion of LP complexity, where we count the number of LP's and their sizes as a complexity measure. This makes sense only when solving LP's dominates other computations such as solving systems of linear equations of sizes of same order. This applies very well to all problems in this Section.

We denote by $\text{LP}(d, m)$ the time necessary to solve any LP with d variables and m inequality constraints: $\max c^T x$ subject to $Ax \leq b$, where A is $m \times d$ rational matrix. We consider $\text{LP}(d, m)$ is an upper bound time measured by big-oh O notation, such as $O(md^3)$ or $O(e^{\sqrt{d \log m}})$. Unlike the usual way to measure the LP complexity using the binary encoding length L of input, we simply ignore L . The main reason is that practically all of implementations of LP algorithms depend hardly on L , but essentially and polynomially on d and m . Further more, we are mostly interested in the case when m is much larger than d and at least as large as $2d$. This practical observation leads to that

Assumption 8.4 *We assume that $\text{LP}(d, m)$ satisfy the following assumptions.*

- (a) $\text{LP}(d, m) = \text{LP}(d + c_1, m + c_2)$ for any constants c_1 and c_2 .
- (b) $\text{LP}(d, m)$ is at least of order md^2 , that is, $\Omega(md^2)$.

The first assumption is based on the fact that LP is solvable in a polynomial time. The second assumption is based on the fact that solving a system of linear inequalities is at least as hard as solving a system of linear equations of the same size (up to constant factor), and the Gaussian elimination has $\Omega(md^2)$ complexity. This second assumption will be used throughout this chapter to argue that the time to solve a linear equality system or to compute a rank of an $m \times d$ matrix is dominated by $\text{LP}(d, m)$.

8.1 Single H-Redundancy Checking

Here we show that Problem 8.1 is linearly equivalent to the linear programming. The one direction is rather obvious, that is, the Single H-redundancy checking can be done by a single LP of the same size.

Proposition 8.5 *Problem 8.1 has No answer if and only if the following LP with $I = [m]$:*

$$\begin{array}{ll} \text{Test}(I, k): & \begin{array}{l} \text{maximize} \quad A_k x \\ \text{subject to} \end{array} \\ & A_i x \leq b_i, \quad \forall i \in I \setminus \{k\} \\ & A_k x \leq b_k + 1 \end{array} \quad (8.2)$$

has an optimal solution whose optimal value is strictly greater than b_k .

Note that the reduction is not only polynomial but linear. Surprisingly, there is a linear reduction from the linear programming (the linear feasibility) to the Single H-redundancy.

Proposition 8.6 *The system $Ax \leq b$ is consistent if and only if a special case of Problem 8.1:*

$$\text{is the inequality } x_0 \leq 0 \text{ redundant in } Ax \leq b \text{ and } x_0 \leq 0 \quad (8.3)$$

has No answer.

Proof. Suppose the $x_0 \leq 0$ is redundant. This is equivalent to the statement that there exists no (x, x_0) such that $Ax \leq b$ and $x_0 > 0$. This in turn is equivalent to the inconsistency of $Ax \leq b$. \blacksquare

We have shown the linear equivalence of the Single H-redundancy checking and the LP. This implies that any redundancy checking algorithm is at least as powerful as an LP algorithm.

In the next section, we will see that removing all H-redundancies can be easier than solving m LP's of size (d, m) that takes time $m \times \text{LP}(d, m)$ if the system is highly redundant.

8.2 H-Redundancy Removal

Here we discuss the problem of removing all redundancies from an H-representation of a polyhedron, i.e., Problem 8.2.

We shall assume that the input of Problem 8.2 is “clean” in the sense that the underlying polyhedron $P = \{x : Ax \leq b\}$ is full-dimensional and no inequality is a positive multiple of another one. This assumption can be met if the preprocessing is done, namely, by embedding the polyhedron in an appropriate subspace. This part will be discussed in Section 8.3.

As we have seen in Section 8.1, removing all redundancies can be done in $m \times \text{LP}(d, m)$ time. Can one do better than this? Here we present an algorithm due to Clarkson [15] which runs much faster than the naive algorithm when the number s of nonredundant inequalities is small relative to m .

Let $Ax \leq b$ be an input system. We assume that a point $z \in \mathbb{Q}^d$ is given satisfying $Az < b$, an interior point of the feasible region $P = \{x : Ax \leq b\}$. At the general stage of the algorithm, we have already detected a row index set I such that the inequality $A_i x \leq b_i$ is nonredundant for $Ax \leq b$, for each $i \in I$. Let j be an row index which is not tested yet, i.e. $j \in [m] \setminus I$. Clarkson's algorithm either detects k th inequality is redundant or finds a row index $j \in [m] \setminus I$ such that $A_j x \leq b_j$ is nonredundant.

```

procedure Clarkson( $A, b, z, I, k$ )
begin
    test whether  $A_k x \leq b_k$  is redundant in  $A_{I \cup \{k\}} x \leq b_{I \cup \{k\}}$ 
    by solving the LP Test( $I \cup \{k\}, k$ ) with optimal solution  $x^*$ 
    if nonredundant then
        (c1)      return (1, RayShoot( $A, b, z, x^* - z$ )) //Returns an essential index
    else
        (c2)      return (0,  $k$ ) //Returns an redundant index
    endif
end

```

Here, the procedure RayShoot(A, b, z, r) returns an index j of a facet-inducing hyperplane $\{x : A_j x = b_j\}$ hit by the ray starting from z along the direction r . It can be easily implemented by analyzing symbolically the ray starting from $z + (\epsilon, \epsilon^2, \dots, \epsilon^d)^T$ along the direction r for sufficiently small $\epsilon > 0$.

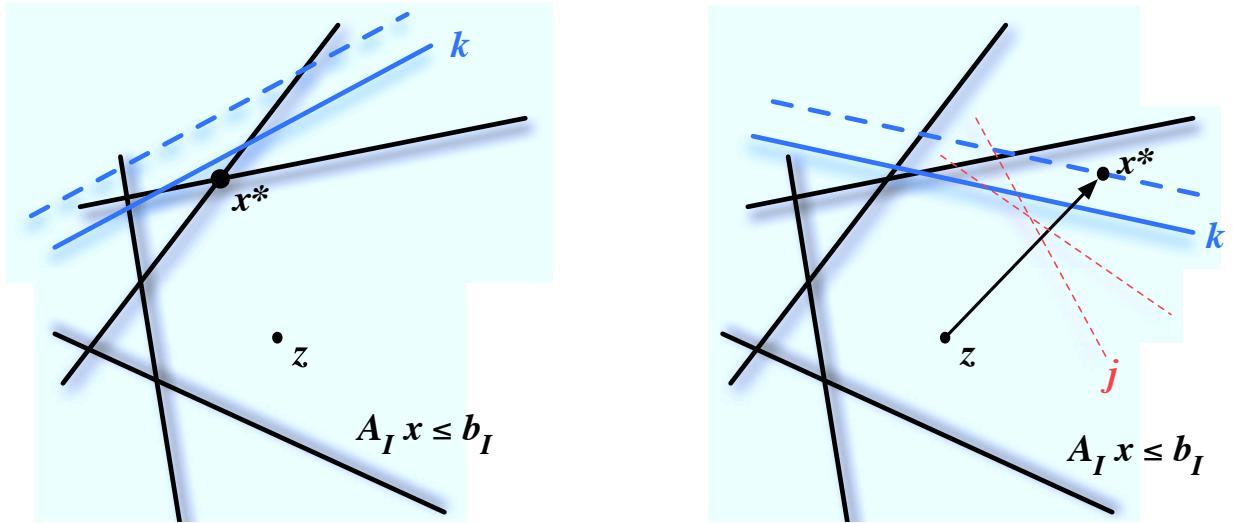


Figure 8.1: Clarkson's Algorithm: Left (k is redundant), Right (An essential j is found)

Exercise 8.1 Write a procedure RayShoot(A, b, z, r) following the specification above. It should be specific enough to be implemented with high level computer languages like C and C++.

Here is the complete algorithm to remove all redundancies. We assume that an interior point z of $P = \{x : Ax \leq b\}$ is given.

```

procedure RedundacyRemovalClarkson( $A, b, z$ )
begin
    set  $I := \emptyset, J := [m]$ 
    repeat
        select an index  $k$  from  $J$ 

```

```

 $(\alpha, j) = \text{Clarkson}(A, b, z, I, k)$ 
 $\text{if } \alpha = 1 \text{ then } I := I \cup \{j\}$  //Increment the essential set  $I$ 
 $J := J \setminus \{j\}$ 
 $\text{until } J = \emptyset$ 
 $\text{return } I$ 
end

```

Theorem 8.7 *The complexity of Clarkson's algorithm to find a minimal equivalent subsystem of $Ax \leq b$ is $m \times \text{LP}(d, s)$ where s is the number of nonredundant constraints in $Ax \leq b$.*

Proof. At each step, Clarkson's algorithm either finds a row index k to be a redundant inequality row index or discovers a new row index $j \neq k$ for which $A_j x \leq b_j$ is essential. Since the size of an LP solved has d variables and at most $s + 1$ constraints, the complexity follows. Note that the complexity of a ray shooting is $O(md)$. Since the number of ray shooting is at most s , the total time $O(smd)$ of ray shooting is dominated by $m \times \text{LP}(d, s)$.

■

8.3 Computing H-Dimension

It is often important to know the dimension of a polyhedron. When a polyhedron is a V-polyhedron with representation, it is very easy to compute its dimension. More precisely, if P is a V-polyhedron for some generator matrices $d \times s$ matrix V and $d \times t$ matrix R , i.e.,

$$P = \{x : x = V\lambda + R\mu, \mathbf{1}^T \lambda = 1, \lambda \geq \mathbf{0}, \mu \geq \mathbf{0}\},$$

then the dimension of P is easily computable, namely by the formula,

$$\dim P = \text{rank} \begin{bmatrix} V & R \\ \mathbf{1}^T & \mathbf{0}^T \end{bmatrix} - 1.$$

However, for an H-polyhedron

$$P = \{x : Ax \leq b\}$$

its dimension is nontrivial to compute. Why nontrivial? It is simply because if one knows the dimension, one can decide whether P is empty or not, that is the linear feasibility problem, equivalent to LP. Then, the next question is how many LP's one has to solve to determine the dimension. Obviously, at least one. It is not hard to see that at most m LP's is sufficient.

In this section, we show that one can compute the dimension by solving at most d LP's. As a byproduct, one also finds a point in the relative interior of P .

The first step is to try to find an interior point of P . If it is successful, the dimension is of course d . One can easily see that the following LP will detect the full-dimensionality:

$$\begin{aligned}
 & \text{maximize} && x_0 \\
 & \text{subject to} && \\
 & && Ax + \mathbf{1}x_0 \leq b, \\
 & && x_0 \leq 1.
 \end{aligned} \tag{8.4}$$

More precisely, we have three cases, depending on the outcome of the LP. Let x^* be an optimal solution and let x_0^* be the optimal value.

Case 1: $x_0^* > 0$. In this case, an optimal solution x^* is an interior point and $\dim P = d$.

Case 2: $x_0^* < 0$. In this case, the polyhedron P is empty and $\dim P = -1$.

Case 3: $x_0^* = 0$. In this case, the polyhedron P is neither full-dimensional nor empty.

In case 3, we must do more computation. For that, we can make use of a dual optimal solution (s^*, t^*) for the dual LP:

$$\begin{aligned} & \text{minimize} && b^T s + t \\ & \text{subject to} && \\ & A^T s &=& \mathbf{0}, \\ & \mathbf{1}^T s + t &=& 1, \\ & s &\geq& \mathbf{0}, \quad t \geq 0. \end{aligned} \tag{8.5}$$

By strong duality, the dual optimal value is zero. This means that s^* cannot be totally zero. Let $I = \{i : s_i^* > 0\}$. By the complementary slackness, at any feasible solution (x, x_0) with $x_0 = 0$ (i.e., at any solution x for $Ax \leq b$), every inequality in $A_I x \leq b_I$ must be tight. We might do even further. By Gaussian elimination, we can recognize all other inequalities in $Ax \leq b$ that are forced to be equalities provided $A_I x = b_I$. Let us merge I with these dependent equality indices, and call it \hat{I} . Now we are ready to solve another LP to find more implicit equalities in the remaining system. For this, let $C := \hat{I}$, and $D := [m] \setminus C$, and set up an LP:

$$\begin{aligned} & \text{maximize} && x_0 \\ & \text{subject to} && \\ & A_C x &=& b_C, \\ & A_D x + \mathbf{1} x_0 &\leq& b_D, \\ & x_0 &\leq& 1. \end{aligned} \tag{8.6}$$

At an optimal solution (x^*, x_0^*) , there are only two cases because x^* cannot be negative this time. When $x_0^* > 0$, the solution x^* is a relative interior point, and the dimension of P is easily computed. It is d minus the maximum number of independent equalities in $A_I x = b_I$. When $x_0^* = 0$, we do essentially the same thing as we did at the very first stage: use the dual optimal solution to recognize implicit equalities in $A_D x \leq b_D$. Then extend them with possible dependent equalities. Another LP should be solved with extended C and its complement D . Since every time an LP is solved, at least one independent implicit equality is found. This shows that at most d LP's will be solved until a relative interior point is found. Thus we have:

Theorem 8.8 *Problem 8.3 can be solved in $d \times \text{LP}(d, m)$ time.*

Exercise 8.2 (Embedding a Polyhedron) Given a point z in the relative interior of $P = \{Ax \leq b\}$, explain a method to embed P to a lower dimensional space so that it is full-dimensional there.

8.4 Reducing the Nonhomogeneous Case to the Homogeneous Case

We define the *homogenization* of a system $Ax \leq b$ as the new system with one extra non-negative variable x_0 ,

$$Ax \leq b x_0 \text{ and } x_0 \geq 0. \quad (8.7)$$

Proposition 8.9 *Let $Ax \leq b$ be a consistent system. An inequality $A_i x \leq b_i$ is redundant in the system if and only if the corresponding inequality $A_i x \leq b_i x_0$ is redundant in the homogenization.*

Exercise 8.3 Prove the proposition above. Show that the assumption that $Ax \leq b$ being consistent is necessary by providing a small example in which $A_i x \leq b_i$ is redundant in $Ax \leq b$ but $A_i x \leq b_i x_0$ is nonredundant in the homogenization.

What we have shown above is that the H-redundancy removal for cones solves the more general problem for polyhedra.

What about for a V-polyhedron? Can we reduce the redundancy removal for V-polyhedra to the one for V-cones? Consider a V-polyhedron with generator pair (V, R) where $V \in \mathbb{Q}^{d \times s}$ and $R \in \mathbb{Q}^{d \times t}$:

$$P_V(V, R) := \{x : x = V\lambda + R\mu, \mathbf{1}^T \lambda = 1, \lambda \geq \mathbf{0}, \mu \geq \mathbf{0}\}.$$

Let v_j denote the j th column of V , and r_k denote the k th column of R . We say a generator v_j (r_k) is *redundant* for $P_V(V, R)$ if removing v_j from V (r_k from R , respectively) does not alter the polyhedron.

Proposition 8.10 *For $V \in \mathbb{Q}^{d \times s}$ and $R \in \mathbb{Q}^{d \times t}$, a generator v_j (r_k) is redundant for $P_V(V, R)$ if and only if the corresponding generator $\begin{bmatrix} v_j \\ 1 \end{bmatrix}$ ($\begin{bmatrix} r_k \\ 0 \end{bmatrix}$, respectively) is redundant in the homogenization*

$$C_V(\hat{R}) := \{x : x = \hat{R}\mu, \mu \geq \mathbf{0}\}, \text{ where } \hat{R} = \begin{bmatrix} V & R \\ \mathbf{1}^T & \mathbf{0}^T \end{bmatrix}.$$

Proof. The proof is straightforward. Left to the reader. ■

Now, we know that both the H-redundancy removal and the V-redundancy removal for cones are as powerful as those for polyhedra. Finally, we have the duality of H- and V-redundancy removals which implies that an algorithm for one problem solves both.

Proposition 8.11 *Let $A \in \mathbb{Q}^{m \times d}$. Then, each inequality $A_i x \leq 0$ is redundant in $Ax \leq \mathbf{0}$ if and only if the corresponding generator A_i^T is redundant for $C_V(A^T)$.*

Proof. Let $A_i x \leq 0$ be redundant in $Ax \leq \mathbf{0}$. This means there exists no x such that $A_i x > 0$ and $A_j x \leq 0$ for all $j \neq i$. By the Farkas Lemma (Exercise 3.4), this is equivalent to the existence of $\mu \geq \mathbf{0}$ such that $A_i^T = \sum_{j \neq i} A_j^T \mu_j$. This is equivalent to saying A_i^T is redundant for $C_V(A^T)$. This completes the proof. ■

Exercise 8.4 (Dimensionality and Linearity) Given a point z in the relative interior of $C_H(A) := \{Ax \leq 0\}$, explain a method to find a basis of the linearity space of $C_V(A^T)$.

9 Polyhedral Representation Conversion

The Minkowski-Weyl Theorem, Theorem 3.9, shows that every convex polyhedron has two representations, an H-representation and a V-representation. The associated problem of computing a (minimal) V-representation from a H-representation or its converse is known as the *representation conversion problem* for polyhedra.

One important characteristic of the representation conversion problem is that the size of output is not easy to measure in terms of the size of input. For example, for a d -cube having $2d$ facets and 2^d vertices, the H-to-V conversion has output whose size is exponential in the input size and the V-to-H conversion has the output size very small relative to the input size.

Given this diversity of output sizes, an ideal algorithm for the conversion problem must be sensitive to the output size, as opposed to optimal for the worst-case output size of a given input size. An algorithm is called *output-polynomial* if its running time is bounded by a polynomial function of both the input size and the output size.

Also, we must take account of the memory footprint. Some algorithms need to store a large amount of data in the memory, while others simply do not store anything except the input data and a few more. We say an algorithm is *compact* if its space is bounded by a polynomial function of the input size only.

One might call an algorithm ideal if it is both compact and output-polynomial. For the representation conversion problem, there is no known output-polynomial algorithm in general. However for the special cases of various nondegeneracy, compact output-polynomial algorithms are known, typically based on the reverse-search paradigm, see Section 9.2.

9.1 Incremental Algorithms

In this section, we present a classical finite algorithm, known as the double description (DD) method [40]. It can be also considered as a constructive proof of Minkowski's Theorem, the implication of (a) \implies (b) in the Minkowski-Weyl Theorem, Theorem 3.10. The algorithm is not output-polynomial as it was shown by Bremner [10]. However, it is extremely useful for certain representation conversion problems, in particular, for highly degenerate inputs.

Suppose that an $m \times d$ matrix A is given, and let $C(A) = \{x : Ax \leq 0\}$. We call any vector $r \in C(A)$ a *ray* of $C(A)$. The DD method is an incremental algorithm to construct a $d \times n$ matrix R such that (A, R) is a DD pair.

Let K be a subset of the row indices $\{1, 2, \dots, m\}$ of A and let A_K denote the submatrix of A consisting of rows indexed by K . Suppose we already found a generating matrix R for $C(A_K)$, or equivalently (A_K, R) is a DD pair. If $A = A_K$, clearly we are done. Otherwise we select any row index i not in K and try to construct a DD pair (A_{K+i}, R') using the information of the DD pair (A_K, R) . Note that $K + i$ is a simplified notation for $K \cup \{i\}$.

Once this basic procedure is described, we have an algorithm to construct a generating matrix R for $C(A)$. This procedure can be easily understood geometrically and the reader is strongly encouraged to draw some simple example in the three dimensional space.

The newly introduced inequality $A_i x \leq 0$ partitions the space \mathbb{R}^d into three parts:

$$\begin{aligned} H_i^+ &= \{x \in \mathbb{R}^d : A_i x > 0\} \\ H_i^0 &= \{x \in \mathbb{R}^d : A_i x = 0\} \\ H_i^- &= \{x \in \mathbb{R}^d : A_i x < 0\}. \end{aligned} \quad (9.1)$$

Let J be the set of column indices of R and let r_j denote the j th column of R . The rays r_j ($j \in J$) are then partitioned into three parts:

$$\begin{aligned} J^+ &= \{j \in J : r_j \in H_i^+\} \\ J^0 &= \{j \in J : r_j \in H_i^0\} \\ J^- &= \{j \in J : r_j \in H_i^-\}. \end{aligned} \quad (9.2)$$

We call the rays indexed by J^+ , J^0 , J^- the *positive*, *zero*, *negative* rays with respect to i , respectively. To construct a matrix R' from R , we generate new $|J^+| \times |J^-|$ rays lying on the i th hyperplane H_i^0 by taking an appropriate positive combination of each positive ray r_j and each negative ray $r_{j'}$ and by discarding all positive rays.

The following lemma ensures that we have a DD pair (A_{K+i}, R') , and provides the key procedure for the most primitive version of the DD method.

Lemma 9.1 (Main Lemma for Double Description Method) *Let (A_K, R) be a DD pair and let i be a row index of A not in K . Then the pair (A_{K+i}, R') is a DD pair, where R' is the $d \times |J'|$ matrix with column vectors r_j ($j \in J'$) defined by*

$$\begin{aligned} J' &= J^- \cup J^0 \cup (J^+ \times J^-), \text{ and} \\ r_{jj'} &= (A_i r_j) r_{j'} - (A_i r_{j'}) r_j \text{ for each } (j, j') \in J^+ \times J^-. \end{aligned}$$

Proof. Let $C = C(A_{K+i})$ and let C' be the cone generated by the matrix R' . We must prove that $C = C'$. By the construction, we have $r_{jj'} \in C$ for all $(j, j') \in J^+ \times J^-$ and $C' \subset C$ is clear.

Let $x \in C$. We shall show that $x \in C'$ and hence $C \subseteq C'$. Since $x \in C$, x is a nonnegative combination of r_j 's over $j \in J$, i.e., there exist $\lambda_j \geq 0$ for $j \in J$ such that

$$x = \sum_{j \in J} \lambda_j r_j. \quad (9.3)$$

If there is no positive λ_j with $j \in J^+$ in the expression above then $x \in C'$. Suppose there is some $k \in J^+$ with $\lambda_k > 0$. Since $x \in C$, we have $A_i x \leq 0$. This together with (9.3) implies that there is at least one $h \in J^-$ with $\lambda_h > 0$. Now by construction, $hk \in J'$ and

$$r_{hk} = (A_i r_h) r_k - (A_i r_k) r_h. \quad (9.4)$$

By subtracting an appropriate positive multiple of (9.4) from (9.3), we obtain an expression of x as a positive combination of some vectors r_j ($j \in J'$) with new coefficients λ_j where the number of positive λ_j 's with $j \in J^+ \cup J^-$ is strictly smaller than in the first expression. As long as there is $j \in J^+$ with positive λ_j , we can apply the same transformation. Thus we must find in a finite number of steps an expression of x without using r_j with $j \in J^+$. This proves $x \in C'$, and hence $C \subseteq C'$. \blacksquare

This algorithm can be used to prove Minkowski's Theorem constructively.

Proof. (of Theorem 3.10) By Lemma 9.1, it is sufficient to show that one can find an initial DD pair (A_K, R) for some K . The trivial case is when $K = \emptyset$ and $C(A_K) = \mathbb{R}^d$. In this case, the set of $2d$ vectors $R = \{e_1, -e_1, e_2, -e_2, \dots, e_d, -e_d\}$ generates the space \mathbb{R}^d by their nonnegative combinations. (Actually, one can find $d + 1$ vectors which positively span \mathbb{R}^d . How?) This completes the proof. \blacksquare

Here we write the DD method in procedural form.

```

procedure DoubleDescriptionMethod( $A$ );
begin
    Obtain any initial DD pair  $(A_K, R)$ 
    while  $K \neq \{1, 2, \dots, m\}$  do
        begin
            Select any index  $i$  from  $\{1, 2, \dots, m\} \setminus K$ 
            Construct a DD pair  $(A_{K+i}, R')$  from  $(A_K, R)$ 
            /* by using Lemma 9.1 */
             $R := R'$     $K := K + i$ ;
        end
        Output  $R$ 
    end.

```

The DD method given here is very primitive, and the straightforward implementation is not quite useful, because the size of J increases very fast and goes beyond any tractable limit. One reason for this is that many (perhaps, most) vectors $r_{jj'}$ the algorithm generates (defined in Lemma 9.1), are unnecessary. To avoid generating redundant generators, we store and update the adjacency of generators. Such a refinement can reduce the size of the output drastically.

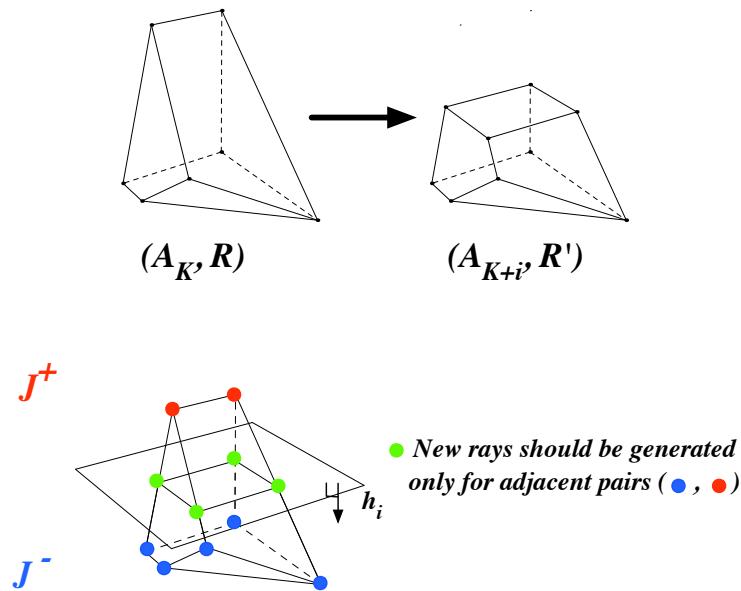


Figure 9.1: The Double Description Algorithm

Figure 9.1 depicts a general step of adding the i th inequality with a refined double description algorithm. The polytopes should be considered as a cut section of 4-dimensional pointed cones with some hyperplane so that each vertex represents one dimensional extreme ray starting from the origin.

Two generators are said to *adjacent* if the common set of active constraints is maximal among all pairs of generators. This means that the line segment connecting a adjacent pair meets the new hyperplane $h_i = \{x : A_i x = 0\}$ at a point lying on a minimal face of the cone $C(A_{K+i})$. It is easy to see that such a point must be in any V-representation.

The double description algorithm at the ideal form not generating any redundant generators is still not easy to analyze. The main problem is that the size of a V-representation of intermediate cone is not easy to estimate. The size also depends heavily on the insertion order of constraints.

Here are somewhat surprising behaviors of the refined double description method with respect to different insertion orders. In the figure below, the input is (the homogenized cone of) a 15-dimensional polytope with 32 facets. The output is a list of 368 vertices. It is important to note that the conversion is **highly degenerate**, meaning that the number of active inequalities at each output vertex is much higher than the dimension.

We consider the five different orderings of the inequalities. The ordering **lexmin** is simply sort the rows of A by lexicographic ordering, comparing the first component first, then the second in case of tie, and the third, etc. The ordering **maxcutoff** (**mincutoff**) is a dynamic ordering in which at each iteration the next inequality is selected to maximize (minimize) the size $|J^+|$. The **lexmin** is a sort of shelling ordering which appears to perform the best among all orderings tested.

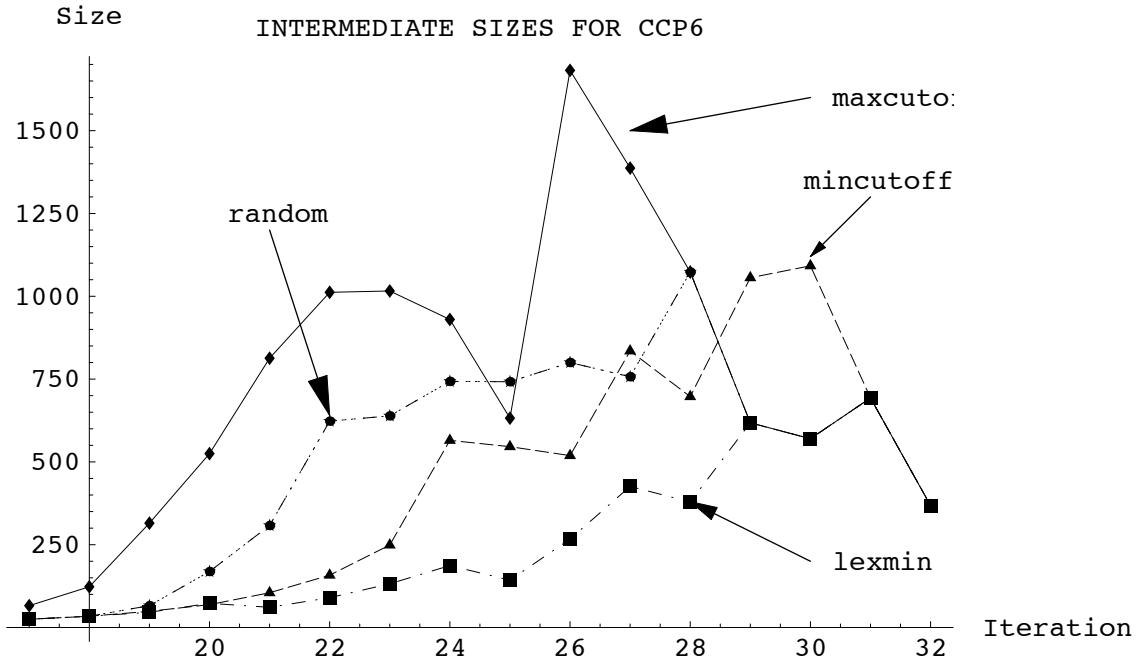


Figure 9.2: Comparison of Intermediate Sizes for a Degenerate Input

A more striking comparison is given below where the input is a 10-dimensional cross polytope with 2^{10} facets. The output is a list of 20 vertices. The highest peak is attained by maxcutoff ordering, following by random and mincutoff. The ordering lexmin is the best among all and the peak intermediate size is less than 30. Note that the graph of lexmin is too low to see it in the figure below.

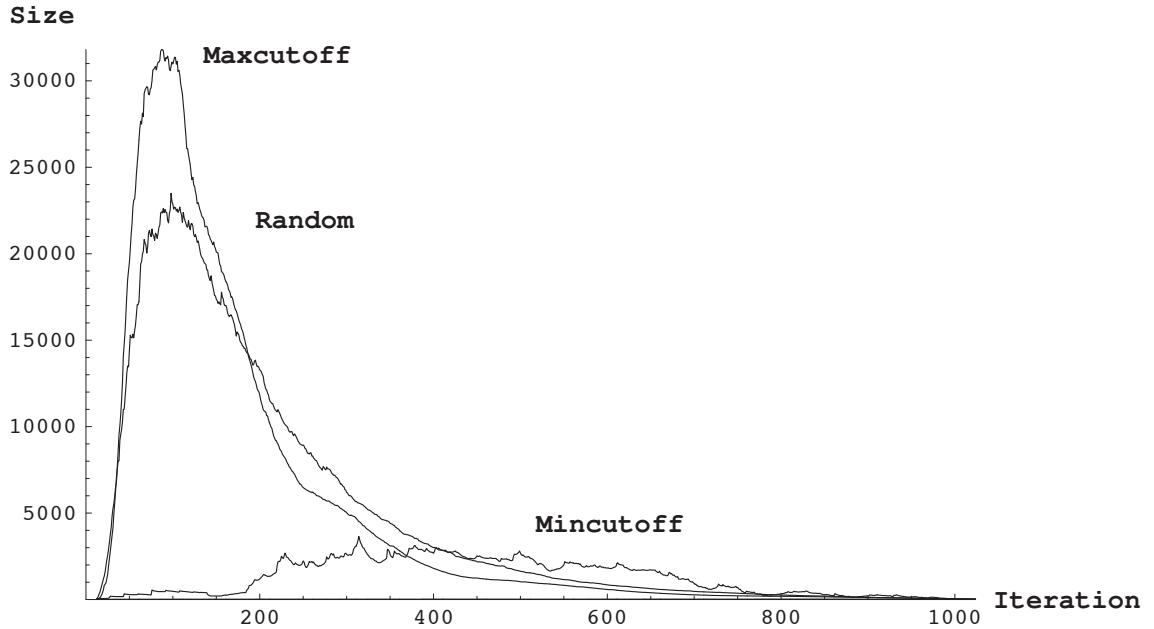
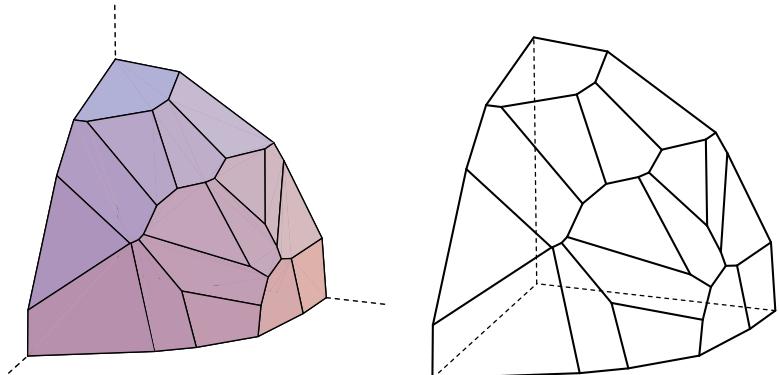


Figure 9.3: Comparison of Intermediate Sizes for a Highly Degenerate Input

9.2 Pivoting Algorithms

One can design pivoting algorithms to visit all vertices of a convex polytope systematically. The idea is quite simple. The graph of a convex polytope is connected, and in fact d -connected if the polytope is d -dimensional, due to Balinski [7]. Thus, one can trace the graph systematically until no new vertex can be found.



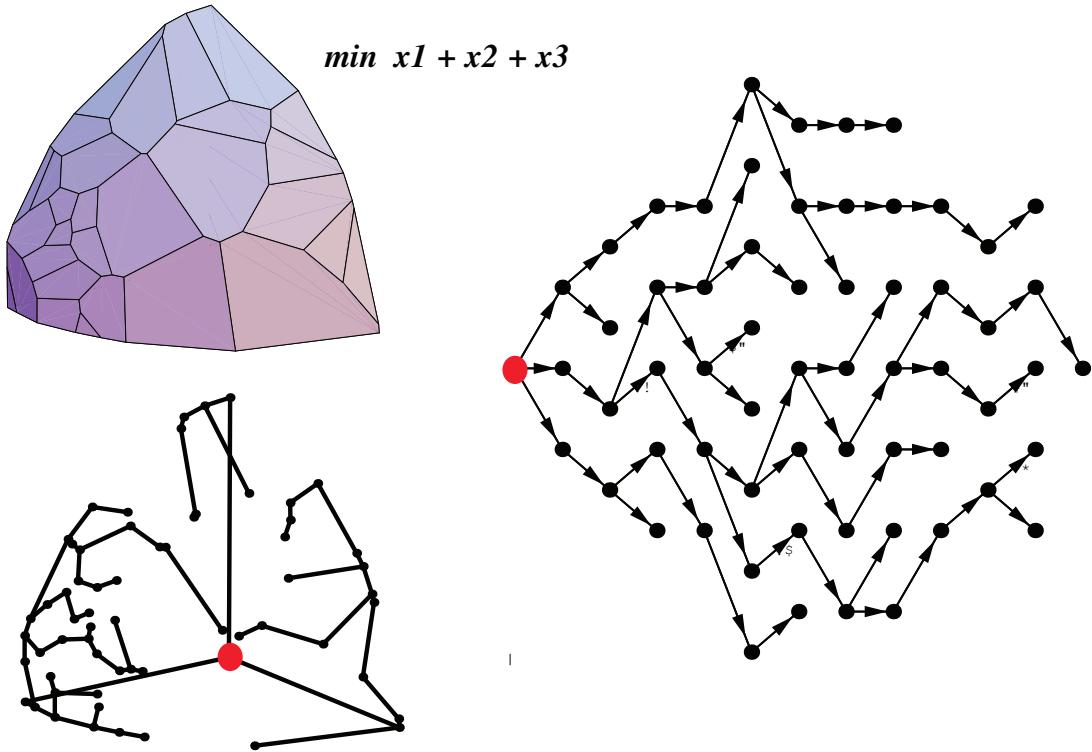
A polytope P

and its graph (1-skeleton)

Historically, there are many pivoting algorithms proposed by Balinski [6], Murty [42], Dyer and Proll [19], etc. The weakness of pivoting algorithms is that when the polytope is degenerate, i.e., non-simple, pivoting may not be able to trace the graph of a polytope efficiently. Typical way to resolve degeneracy is a symbolic perturbation of constraints which may create an exponentially large number of new extreme points. Under the nondegeneracy assumption that the number of active constraints at each vertex is exactly d , the algorithm due to Dyer and Proll [19] is an output-polynomial algorithm. Yet, it must store all visited vertices in memory and thus is not a compact algorithm.

In this section, we present a compact output-polynomial algorithm for the nondegenerate case, based on the reverse search technique due to Avis and Fukuda.

The main idea is to reverse the simplex method from the optimal vertex in all possible ways. Here the objective function is set to any generic one so that the optimal vertex is unique and no edge of the polytope is parallel to an objective contour.



Also, another important thing is to make sure that the simplex algorithm is finite and selects a next pivot uniquely at each vertex. This can be achieved, for example, by employing the minimum index rule (Bland's rule). Under these, the edges used by the refined simplex method form a directed spanning tree of the graph $G(P)$ of a polytope P rooted at the optimal vertex. We will see that the resulting algorithm enumerates all f_0 vertices in time $O(mdf_0)$ and $O(md)$ -space under nondegeneracy when the input H-polytope is given by m inequalities in d variables. Thus, it is compact and output-polynomial.

For a formal description, let us define two functions. A finite local search f for a graph $G = (V, E)$ with a special node $s \in V$ is a function: $V \setminus \{s\} \rightarrow V$ satisfying

(L1) $\{v, f(v)\} \in E$ for each $v \in V \setminus \{s\}$, and

(L2) for each $v \in V \setminus \{s\}$, $\exists k > 0$ such that $f^k(v) = s$.

For example, let $P = \{x \in \mathbb{R}^d : A x \leq b\}$ be a simple polytope, and $c^T x$ be any generic linear objective function. Let V be the set of all vertices of P , s the unique optimal, and $f(v)$ be the vertex adjacent to v selected by the simplex method which selects a pivot uniquely if v is not the optimal vertex.

A *adjacency oracle* or simply *A-oracle* Adj for a graph $G = (V, E)$ is a function (where δ a upper bound for the maximum degree of G) satisfying:

- (i) for each vertex v and each number k with $1 \leq k \leq \delta$ the oracle returns $\text{Adj}(v, k)$, a vertex adjacent to v or extraneous *null* (*null*),
- (ii) if $\text{Adj}(v, k) = \text{Adj}(v, k') \neq 0$ for some $v \in V$, k and k' , then $k = k'$,
- (iii) for each vertex v , $\{\text{Adj}(v, k) : \text{Adj}(v, k) \neq 0, 1 \leq k \leq \delta\}$ is exactly the set of vertices adjacent to v .

For example, when $P = \{x \in \mathbb{R}^d : A x \leq b\}$ is a simple polytope, let V be the set of all vertices of P , δ be the number of nonbasic variables and $\text{Adj}(v, k)$ be the vertex adjacent to v obtained by pivoting on the k th nonbasic variable at v .

Now we are ready to describe a general reverse search algorithm to generate all vertices of the underlying graph G assuming that the two functions f and Adj are given.

```

procedure ReverseSearch( $\text{Adj}, \delta, s, f$ );
   $v := s; j := 0$ ; (*  $j$ : neighbor counter *)
  repeat
    while  $j < \delta$  do
       $j := j + 1$ ;
      (r1)  $next := \text{Adj}(v, j)$ ;
      if  $next \neq \text{null}$  then
        (r2) if  $f(next) = v$  then (* reverse traverse *)
           $v := next; j := 0$ 
        endif
      endif
    endwhile;
    if  $v \neq s$  then (* forward traverse *)
      (f1)  $u := v; v := f(v)$ ;
      (f2)  $j := 0; \text{repeat } j := j + 1 \text{ until } \text{Adj}(v, j) = u$  (* restore  $j$  *)
    endif
  until  $v = s$  and  $j = \delta$ 

```

We can evaluate the complexity of reverse search above as follows. Below we denote by $t(f)$ and $t(\text{Adj})$ the time to evaluate the functions f and Adj , respectively.

Theorem 9.2 Suppose that a local search (G, s, f) is given by an *A-oracle*. Then the time complexity of *ReverseSearch* is $O(\delta t(\text{Adj})|V| + t(f)|E|)$.

Proof. It is easy to see that the time complexity is determined by the total time spent to execute the four lines (r1), (r2), (f1) and (f2). The first line (r1) is executed at most δ times for each vertex, and the total time spent for (r1) is $O(\delta t(\text{Adj})|V|)$. The line (r2) is executed as many times as the degree $\deg(v)$ for each vertex v , and thus the total time for (r2) is $O(t(f)|E|)$. The third line (f1) is executed for each vertex v in $V \setminus \{s\}$, and hence the total time for (f1) is $O(t(f)(|V| - |S|))$. Similarly, the total time for (f2) is $O(\delta t(\text{Adj})(|V|))$. Since $|V| \leq |E|$, by adding up the four time complexities above, we have the claimed result. \blacksquare

Corollary 9.3 Suppose that a local search (G, s, f) is given by an A-oracle. Then the time complexity of ReverseSearch is $O(\delta(t(\text{Adj}) + t(f))|V|)$. In particular, if δ , $t(f)$ and $t(\text{Adj})$ are independent of the number $|V|$ of vertices in G , then the time complexity is linear in the output size $|V|$.

Proof. The claim follows immediately from Theorem 9.2 and the fact that $2|E| \leq \delta|V|$. \blacksquare

One can improve the complexity of reverse search algorithms by exploiting special structures. We give the best known complexity of reverse search for the representation conversion for convex polytopes without proof, see [4] for details.

Theorem 9.4 There is an implementation of reverse search algorithm to enumerate all vertices of a nondegenerate H -polytope $P = \{x : Ax \leq b\}$ in time $O(mdf_0)$ and space $O(md)$, where $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$ and f_0 is the number of vertices of P . In particular, it is a compact output-polynomial algorithm for nondegenerate inputs.

There are many applications of reverse search in geometry and combinatorics, see [5].

Finally, what is the difference between reverse search and depth-first search? The quick answer is: reverse search is a memory-free search while depth-first search must store all nodes visited so far to distinguish those vertices from the rest. In other words, reverse search is depth-first search applied to a unique spanning tree of the graph defined by local search function f .

9.3 Pivoting Algorithm vs Incremental Algorithm

- Pivoting algorithms, in particular the reverse search algorithm (lrs, lrslib [2]), work well for high dimensional cases.
- Incremental algorithms work well for low (say, up to 12) dimensional cases and highly degenerate cases. For example, the codes cdd/cddlib [22] and porta [13] are implemented for highly degenerate cases and the code qhull [8] for low (up to 10) dimensional cases.
- The reverse search algorithm seems to be the only method that scales very efficiently in massively parallel environment.
- Various comparisons of representation conversion algorithms and implementations can be found in the excellent article [3] by Avis, Bremner and Seidel.

10 Hyperplane Arrangements and Point Configurations

In Sections 5, 6 and 7, we studied the combinatorial structure of convex polytopes. In this section, we look at not only polytopes but also the dissection of the whole space by a set of hyperplanes which induces a polyhedral complex. Formally, it is known as an arrangement of hyperplanes and its dual structure is known as a configuration of points or vectors.

10.1 Cake Cutting

An intuitively appealing way to study the dissection of the plane by a set of lines is through cake cutting. Just consider a round cake from above (i.e, a 2-dimensional disk), and try to cut it by a knife a few times. With m straight cuts, how many pieces can one produce? Of course, it depends on cut intersection patterns, as one can see in Figure 10.1.

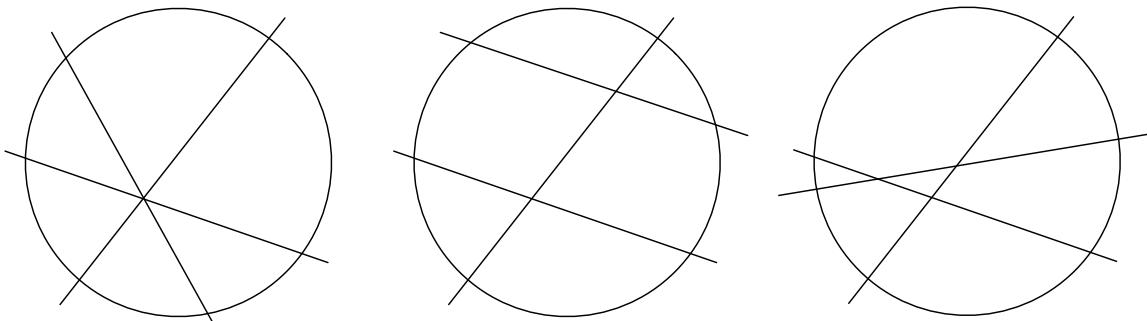
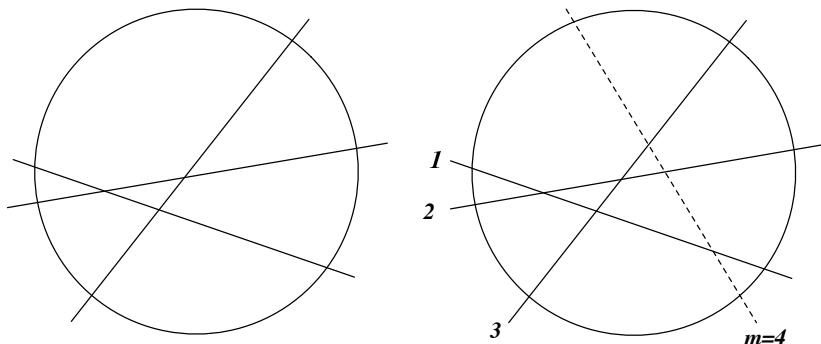


Figure 10.1: Cake Cutting Problem

Let us denote by $p_2(m)$ the maximum number of pieces one can produce by m cuts in 2D. Clearly, $p_2(0) = 1$ and $p_2(1) = 2$. It is not hard to give an explicit formula for this by looking at a simple recursive formula. We can easily see that if the m th cut intersects with the previous $m - 1$ cuts at distinct points (in the interior of the cake), then it generates **additional m pieces**. It is obvious that this is an upper bound of the number of pieces one can generate.



Is this upper bound attainable? We argue that this is always attainable by placing cuts properly. A 2D cake cutting with m cuts is defined to be *nondegenerate* if any two distinct cuts intersect in the interior of the cake and no three distinct cuts have a common intersection. For any m , nondegenerate cuts exist. Just place m cuts so that no two cuts are parallel and no three cuts intersect. If some two cuts do not intersect in the interior of the cake, just dilate the cake (centered at the origin) without changing the cut placements. If the dilation is large enough, all intersections of the lines will be placed inside the cake. Then, shrink the whole space so that the cake becomes the original size.

This observation leads to a simple recursion:

$$p_2(m) = p_2(m-1) + m \quad (10.1)$$

which implies

$$p_2(m) = p_2(0) + 1 + 2 + \cdots + m = 1 + \frac{(m+1)m}{2}. \quad (10.2)$$

Now we go up to one higher dimension. The cake looks like a 3-dimensional ball, and we try to cut out the largest number of pieces with m cuts. We now imagine how a “watermelon” can be dissected into pieces by a knife into pieces, see Figure 10.2. (A cut does not go through the center although the figure shows such a case.)

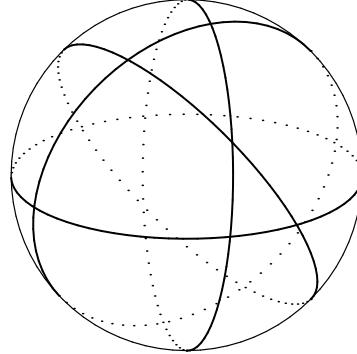


Figure 10.2: 3-Dimensional Cake Cutting Problem

Let us denote by $p_3(m)$ the maximum number of pieces one can produce by m cuts in 3 dimension. Can one write a simple recursive formula for $p_3(m)$? Yes, it is possible, once one notices that the cut section at m th cut in the 3D cake could look like a 2D cake cutting at $(m-1)$ st step, as long as the m th cut intersects with the previous $(m-1)$ cuts at distinct lines. A key observation is that the number of 2D pieces at the cut section is exactly the increment of the number of pieces by m th cut. Thus, when the m th cut section

is nondegenerate, the increment is largest and thus the observation leads to the recursion

$$p_3(m) = p_3(m-1) + p_2(m-1) \quad (10.3)$$

$$\begin{aligned} &= p_3(m-1) + 1 + \frac{m(m-1)}{2} \\ &= p_3(0) + m + \sum_{i=1}^m \frac{i(i-1)}{2} \\ &= 1 + m + \frac{1}{2} \left(\sum_{i=1}^m i^2 - \sum_{i=1}^m i \right) \\ &= 1 + m + \frac{1}{6}(m+1) \left(m + \frac{1}{2} \right) m - \frac{1}{4}m(m+1). \end{aligned} \quad (10.4)$$

$$(\text{Recall the identity: } \sum_{i=1}^m i^2 = \frac{1}{6}(m+1)(2m+1)m)$$

We have two explicit formulas, one for 2D (10.2) and the other for 3D (10.4). Can we guess a general formula for $p_d(m)$? Well, not quite easy to guess from what we have. But, it is much easier once we rewrite the two equations in the following form:

$$\begin{aligned} p_2(m) &= \binom{m}{0} + \binom{m}{1} + \binom{m}{2} \\ p_3(m) &= \binom{m}{0} + \binom{m}{1} + \binom{m}{2} + \binom{m}{3}. \end{aligned} \quad (10.5)$$

Exercise 10.1 Verify the correctness of the equations (10.5).

Now, we are ready to prove the general cake cutting theorem.

Theorem 10.1 *The number $p_d(m)$ of the maximum number of pieces dissected from the d -dimensional ball by m (hyperplane) cuts is given by*

$$p_d(m) = \sum_{i=0}^d \binom{m}{i}. \quad (10.6)$$

Proof. We prove the correctness of the formula and the fact that the value is attained by any nondegenerate cake cut, by induction on d . Here, we say a d -dimensional cake cutting of a d -ball (cake) with m cuts is defined to be *nondegenerate* if any d distinct cuts intersect in the interior of the cake and no $(d+1)$ distinct cuts have a common intersection. The formula is correct for $d = 2$ and attained by any nondegenerate cutting. Consider any unknown case d assuming that the formula is correct for any smaller dimension. First of all, $m = 0$, the formula $p_d(m)$ is correct that is 1. Here we use second induction on m . Consider any unknown case m assuming that the formula is correct for any smaller values of m . By extending the recursion (10.3), we have

$$p_d(m) = p_d(m-1) + p_{d-1}(m-1). \quad (10.7)$$

By induction hypothesis, one can apply the formula to the RHS and we have:

$$\begin{aligned} p_d(m) &= \binom{m-1}{0} + \binom{m-1}{1} + \cdots + \binom{m-1}{d} \\ &\quad + \binom{m-1}{0} + \cdots + \binom{m-1}{d-1}. \end{aligned}$$

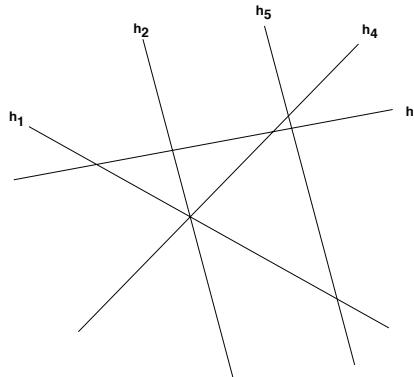
Finally, since $\binom{m-1}{-1} = 0$, the last equation above leads to

$$\begin{aligned} p_d(m) &= \sum_{k=0}^d \left(\binom{m-1}{k} + \binom{m-1}{k-1} \right) \\ &= \sum_{k=0}^d \binom{m}{k}. \end{aligned} \tag{10.8}$$

This completes the proof. ■

10.2 Arrangements of Hyperplanes and Zonotopes

Cake cutting is a less formal way of presenting the mathematical notion of arrangements of hyperplanes in \mathbb{R}^d . A finite family $\mathcal{A} = \{h_i : i = 1, 2, \dots, m\}$ of hyperplanes in \mathbb{R}^d is called an *arrangement of hyperplanes*.

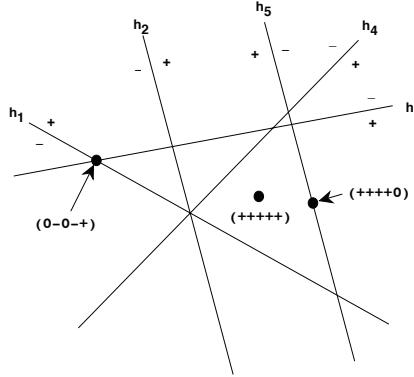


We are mostly interested in combinatorial structures underlying hyperplane arrangements. For this, it is convenient to define the partition of the space \mathbb{R}^d into three sets:

$$h_i^+ = \{x : A_i x < b_i\}, \tag{10.9}$$

$$h_i^0 = \{x : A_i x = b_i\}, \tag{10.10}$$

$$h_i^- = \{x : A_i x > b_i\}. \tag{10.11}$$

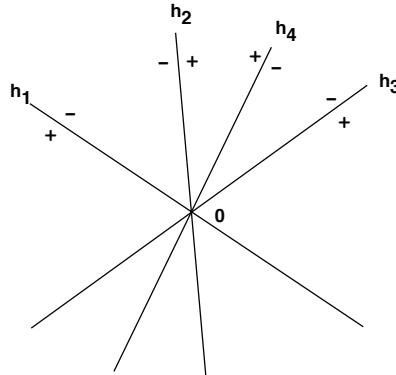


There is a natural way to associate each point x in the space \mathbb{R}^d with the sign vector $\sigma(x) \in \{-, 0, +\}^m$ defined by:

$$\sigma(x)_i = \begin{cases} + & \text{if } x \in h_i^+ \\ 0 & \text{if } x \in h_i^0 \\ - & \text{if } x \in h_i^- \end{cases} \quad i \in E.$$

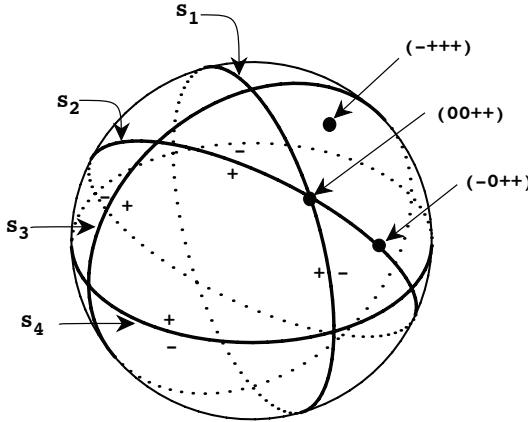
The set of points with a given sign vector is a relatively open polyhedron, is called a *open face* of the arrangement, and its topological closure is called a *face* of the arrangement. The full dimensional faces are called the *cells* or *regions* of the arrangement. The set of all faces forms a polyhedral complex, called the *complex of the arrangement*. One can represent the facial incidence in the complex by a binary relation among sign vectors. For two sign vectors $X, Y \in \{-, 0, +\}^m$, we say X *conforms to* Y (denoted as $X \preccurlyeq Y$) if $i \in [m]$ and $X_i \neq 0$ implies $X_i = Y_i$. The poset $\sigma(\mathbb{R}^d) := \{\sigma(x) : x \in \mathbb{R}^d\}$ ordered by conformal relation is a combinatorial representation of the complex. This poset is the *face poset $\mathcal{F}(\mathcal{A})$ of the arrangement \mathcal{A}* .

The poset $\mathcal{F}(\mathcal{A})$ behaves nicely if all the hyperplanes contains the origin. An arrangement of hyperplanes in which all its hyperplanes contain the origin 0 is called a *central arrangement of hyperplanes*.



For example, $\mathcal{F}(\mathcal{A})$ contains the zero vector $\mathbf{0}$ which is the unique smallest element. Also, it is symmetric with respect to the origin: if a sign vector X is in $\mathcal{F}(\mathcal{A})$, its negative $-X$ is in $\mathcal{F}(\mathcal{A})$. By adding the artificial greatest element $\mathbf{1}$ of all 1's to $\mathcal{F}(\mathcal{A})$, we obtain what we call the face lattice $\hat{\mathcal{F}}(\mathcal{A})$ of the central arrangement. We will see this lattice is isomorphic to the lattice of a very special polytope.

Let \mathcal{A} be a central arrangement of hyperplanes represented by a matrix A , i.e., $h_i = \{x : A_i x = 0\}, \forall i = 1, \dots, m$. Geometrically, it is convenient to look at the cut section of the arrangement with the unit $(d-1)$ -sphere $S^{d-1} := \{x \in \mathbb{R}^d : \|x\| = 1\}$, where each hyperplane becomes a $(d-2)$ -sphere $s_i := h_i \cap S^{d-1}$. Thus, the cut section is an arrangement of $(d-2)$ -spheres in the unit sphere S^{d-1} . The complex of the arrangement is essentially represented in the sphere arrangement, namely, $\sigma(\mathbb{R}^d) = \sigma(S^d) \cup \{\mathbf{0}\}$.



Consider the following H-polyhedron given by 2^m inequalities:

$$P_A = \{x : y^T A x \leq 1, \forall y \in \{-1, +1\}^m\}.$$

Theorem 10.2 *Let A be a column full rank matrix representing a central arrangement \mathcal{A} . Then P_A is a polytope, and the face lattice $\hat{\mathcal{F}}(\mathcal{A})$ of \mathcal{A} is isomorphic to the face lattice of the polytope P_A .*



The central arrangement \mathcal{A}

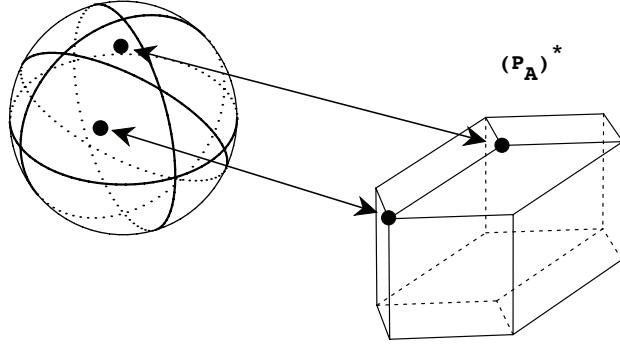
and

the polytope P_A

The polar of the polytope P_A is a very special polytope. In fact, it is a zonotope.

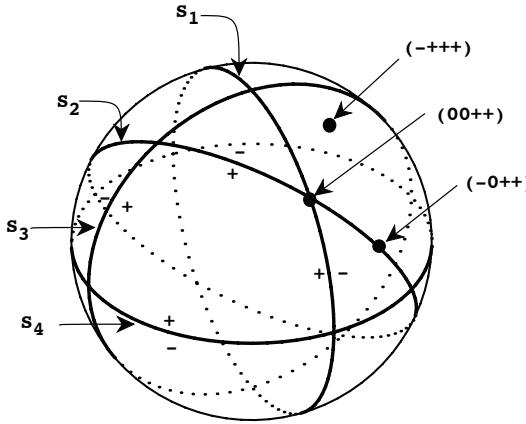
$$\begin{aligned} (P_A)^* &= \text{conv}\{y^T A \in \mathbb{R}^d : y \in \{-1, +1\}^m\} \\ &= \{y^T A \in \mathbb{R}^d : y \in [-1, +1]^m\} \\ &= L_1 + L_2 + \dots + L_m, \end{aligned}$$

where each *generator* L_i is the line segment $[-A_i, A_i]$.



10.3 Face Counting Formulas for Arrangements and Zonotopes

We denote by $f_k(\mathcal{A})$ the number of k -dimensional faces of an arrangement \mathcal{A} of hyperplanes in \mathbb{R}^d . We assume all arrangements are central and thus can be seen as a sphere arrangement in S^{d-1} .



With the sphere arrangement setting, it is not hard to relate any central arrangement of m hyperplanes in \mathbb{R}^d to a cake cutting. Let s_m^0 be the last sphere in the arrangement. It is the boundary of two hemispheres $s_m^+ := h_m^+ \cap S^{d-1}$ and $s_m^- := h_m^- \cap S^{d-1}$. The arrangement of spheres restricted to one of the hemispheres is combinatorially equivalent to the cake cutting of a $d-1$ -dimensional ball by $m-1$ cuts. This observation together with Theorem 10.1 implies the following theorem.

Theorem 10.3 (Upper Bound Theorem for Arrangements) *For any central arrangement \mathcal{A} of m hyperplanes in \mathbb{R}^d ,*

$$f_d(\mathcal{A}) \leq 2 \sum_{i=0}^{d-1} \binom{m-1}{i} \quad \text{and} \quad f_1(\mathcal{A}) \leq 2 \binom{m}{d-1}.$$

Note that if one restrict the arrangement to the unit sphere, the LHS expressions represent $f_{d-1}(\mathcal{A} \cap S^{d-1})$ and $f_0(\mathcal{A} \cap S^{d-1})$.

Using the duality of arrangements and zonotopes, Theorem 10.3 implies the upper bound theorem for zonotopes.

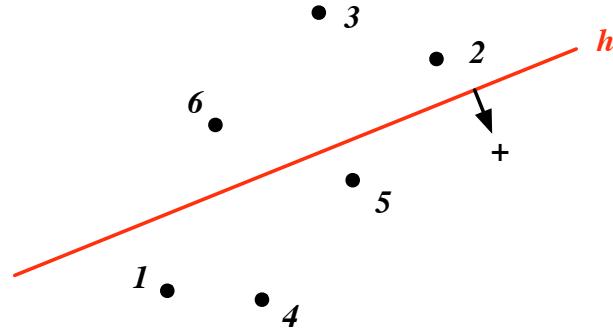
Theorem 10.4 (Upper Bound Theorem for Zonotopes) *Let P be a d -dimensional zonotope given by m generators ($m \geq d$). Then,*

$$f_0(P) \leq 2 \sum_{i=0}^{d-1} \binom{m-1}{i} \quad \text{and} \quad f_{d-1}(P) \leq 2 \binom{m}{d-1}.$$

For fixed d , both $f_{d-1}(P)$ and $f_0(P)$ are $O(m^{d-1})$.

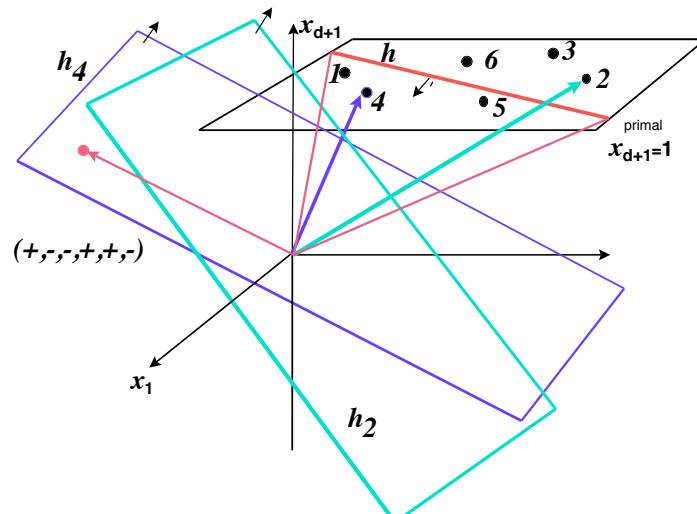
10.4 A Point Configuration and the Associated Arrangement

A *point configuration* is a set $P = \{p_1, p_2, \dots, p_n\}$ of points in \mathbb{R}^d . The relative locations of the points with respect to an arbitrary hyperplane represent the underlying combinatorial structure.



Let $\hat{p}_i = \begin{bmatrix} p_i \\ 1 \end{bmatrix}$ be the lifted points in \mathbb{R}^{d+1} , and the hyperplanes $h_i = \{x : \hat{p}_i^T x = 0\}$. The resulting arrangement $\mathcal{A} = \{h_1, \dots, h_n\}$ in \mathbb{R}^{d+1} encodes the combinatorial structure of P nicely.

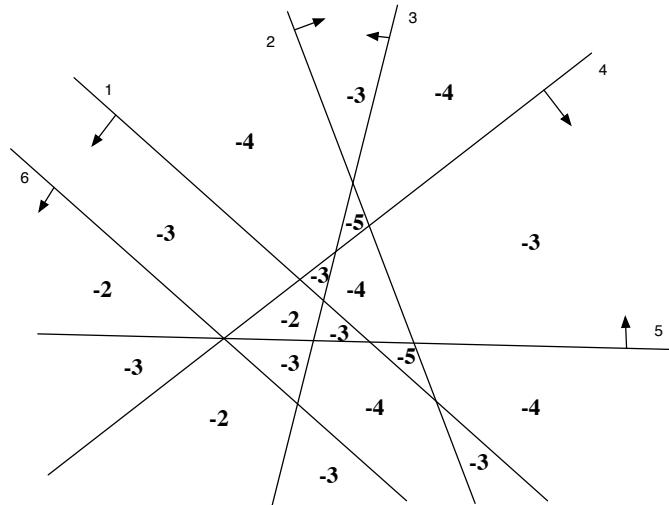
A open halfspace h^+ is represented by the sign vector $X \in \{+, -, 0\}^n$ of a region in the dual hyperplane arrangement with $j \in X^+$ iff $p_j \in h^+$.



The partition $(\{1, 4, 5\}, \{2, 3, 6\})$ by the hyperplane h corresponds to the region $(+, -, -, +, +, -)$.

10.4.1 Application: Largest Feasible Subsystem

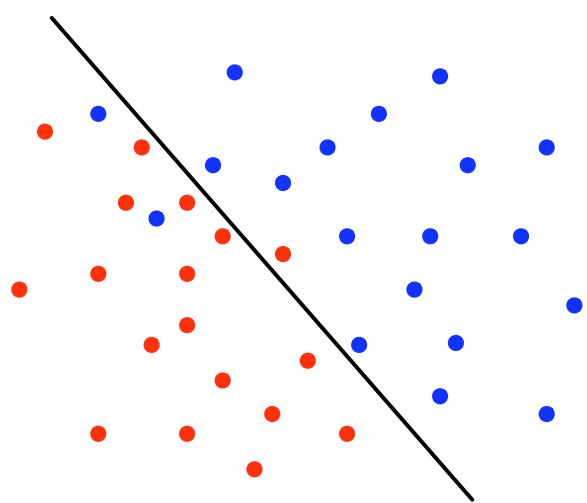
Given an inconsistent linear inequality system $Ax < b$, find a subsystem that is consistent and largest possible. In other words, try to remove as few inequalities as possible to make it feasible.



This problem is known to be NP-hard. One must rely on some kind of enumeration or approximation algorithms to solve this.

10.4.2 Applications: Best Separation of Points by a Hyperplane

Given two blue and red sets of points in \mathbb{R}^d , find a (separation) hyperplane which is best possible, i.e. the number of misclassified points is minimized.



This problem is NP-hard, and in fact, one can reduce this to the largest feasible subsystem problem. The number of separations represents the underlying complexity of enumeration.

11 Computing with Arrangements and Zonotopes

As we learned in the previous section that central arrangements of hyperplanes and zonotopes are essentially the same object mathematically. More specifically, if $\mathcal{A}(A)$ is a central arrangement with an $m \times d$ representation matrix A , then its face lattice $\hat{\mathcal{A}}$ is anti-isomorphic to the zonotope $Z(A)$ generated by the line segments $L_i = [-A_i, A_i]$, $j \in [m]$.

This duality implies that one can translate an algorithm for arrangements to an algorithm for zonotopes. In particular, the following pairs of problems with input matrix A given are equivalent.

Problem 11.1 Cell Enumeration for Arrangements/Vertex Enumeration for Zonotopes

- (a) Generating all cells of $\mathcal{A}(A)$.
- (b) Generating all vertices of $Z(A)$.

Problem 11.2 Vertex Enumeration for Arrangements/Facet Enumeration for Zonotopes

- (a) Generating all 1-faces (rays) of $\mathcal{A}(A)$.
- (b) Generating all facets of $Z(A)$.

Problem 11.3 Face Enumeration for Arrangements/Face Enumeration for Zonotopes

- (a) Generating all faces of $\mathcal{A}(A)$.
- (b) Generating all faces of $Z(A)$.

There is a compact output-polynomial algorithm [5] due to Avis and Fukuda for Problem 11.1. Also, there is a worst-case optimal algorithm [20] due to Edelsbrunner, O'Rourke and Seidel for Problem 11.1.

There is a output-polynomial algorithm [49] due to Seymour for Problem 11.2. No compact output-polynomial algorithm is known for Problem 11.2. When input is nondegenerate, Problem 11.2 has a trivial algorithm which is compact and output-polynomial, just go through all $\binom{m}{d-1}$ combinations. This suggests that when input is only “slightly” degenerate, the naive algorithm might be practical.

The paper [27] shows that there is an output-polynomial algorithm to generate all faces of $\mathcal{A}(A)$ from the list of cells. This means that together with the compact output-polynomial algorithm [5] for Problem 11.1, Problem 11.3 can be solved by an output-polynomial algorithm.

11.1 Cell Generation for Arrangements

Here we present the reverse search algorithm [5] which is the only compact output-polynomial algorithm for generating all cells of an arrangement. By duality, this is a compact output-polynomial algorithm for enumerating all vertices of a zonotope.

Let \mathcal{A} be an arrangement of distinct hyperplanes $\{h_i : i \in [m]\}$ in \mathbb{R}^d , where each hyperplane is given by a linear equality $h_i = \{x : A_i x = b_i\}$. The two sides of h_i are

$h_i^+ = \{x : A_i x \geq b_i\}$ and $h_i^- = \{x : A_i x \leq b_i\}$. For each $x \in \mathbb{R}^d$, the sign vector $\sigma(x)$ of x is the vector in $\{-, 0, +\}^m$ defined by

$$\sigma(x)_i = \begin{cases} - & \text{if } x \in h_i^- \\ 0 & \text{if } x \in h_i \\ + & \text{if } x \in h_i^+ \end{cases} \quad (i \in [m]).$$

Let V_{CELL} be the set of sign vectors of points in \mathbb{R}^d whose nonzero support is $[m]$. We can identify each vector c in V_{CELL} with the open cell (open d -face) of the arrangement defined by $\{x : \sigma(x) = c\}$. For two cells c and c' , let $\text{sep}(c, c')$ be the set of separators of c and c' , that is, the set of elements i in $[m]$ such that c_i and c'_i have opposite signs. We say that two cells c and c' are *adjacent* in G_{CELL} if they differ in only one component, or equivalently, $|\text{sep}(c, c')| = 1$. The following lemma is important.

Lemma 11.4 *For any two distinct cells c and c' in V_{CELL} , there exists a cell c'' which is adjacent to c and $\text{sep}(c, c'') \subset \text{sep}(c, c')$.*

Proof. Let c and c' be two distinct cells, and let x (x') be a point in c (in c' , respectively) in general position. Moving from x toward x' on the line segment $[x, x']$, we encounter the sequence of cells: $c_0 = c, c_1, c_2, \dots, c_k = c'$, and we can easily verify that c_1 is adjacent to c and $\text{sep}(c, c_1) \subset \text{sep}(c, c')$. \blacksquare

Let us assume that V contains the cell c^* of all +'s. Lemma 11.4 implies that for each cell c different from c^* , there is a cell c'' which is adjacent to c and $\text{sep}(c^*, c'') \subset \text{sep}(c^*, c)$. Let us define $f_{CELL}(c)$ as such c'' that is lexico-largest (i.e., the unique element in $\text{sep}(c, c'')$ is smallest possible). Then, $(G_{CELL}, S_{CELL}, f_{CELL})$ is a finite local search with $S_{CELL} = \{c^*\}$.

Figure 11.1 describes the trace of the local search on a small example with $d = 2$ and $m = 4$.

By reversing this local search, we obtain an algorithm to list all cells in an arrangement. There are a few things to be explained for an implementation. First, we assumed that the cell c^* of all +'s is given, but we can pick up any cell c in the arrangement, and consider it as the cell of all +'s since replacing some equality $A_i x = b_i$ by $-A_i x = -b_i$ does not essentially change the arrangement. Note that one can obtain an initial cell by picking up any random point in \mathbb{R}^d and perturbing it if it lies on some hyperplanes.

Now, how can we realize $\text{ReverseSearch}(\text{Adj}_{CELL}, \delta_{CELL}, S_{CELL}, f_{CELL})$ in an efficient way? First we can set $\delta_{CELL} = m$ and $S_{CELL} = \{c^*\}$. For any cell $c \in V_{CELL}$ and $k \in M$, the function $\text{Adj}_{CELL}(c, k)$ can be realized via solving an LP of the form

$$\begin{aligned} & \text{minimize (maximize)} && y_k \\ & \text{subject to} && y = Ax - b, \\ & && y_i \geq 0 \text{ for all } i \neq k \text{ with } c_i = +, \\ & && y_i \leq 0 \text{ for all } i \neq k \text{ with } c_i = -, \end{aligned} \tag{11.1}$$

where minimization (maximization) is chosen when $c_k = +$ ($c_k = -$, respectively). The function returns the adjacent cell c' with $\text{sep}(c, c') = \{k\}$ if and only if LP (11.1) has a feasible solution with negative (positive) objective value. The time $t(\text{Adj}_{CELL})$ depends on how an LP with d variables and $m - 1$ inequalities is solved. We denote this as a function $\text{LP}(d, m)$ of m and d , as we used this notation in Section 8.

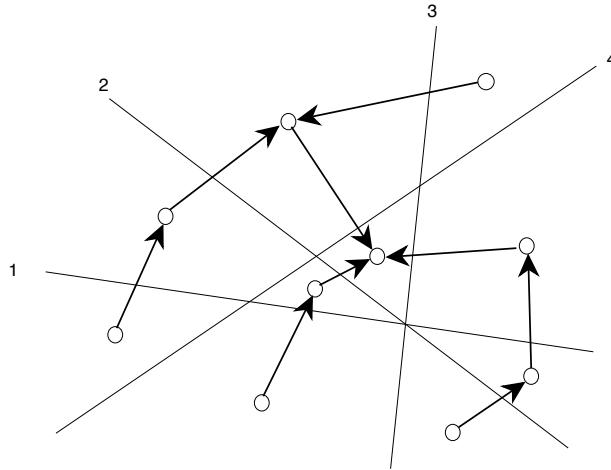
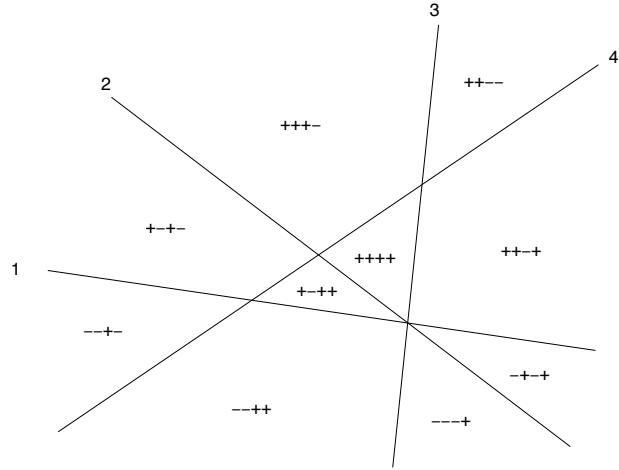


Figure 11.1: An arrangement of hyperplanes and the trace of f_{CELL}

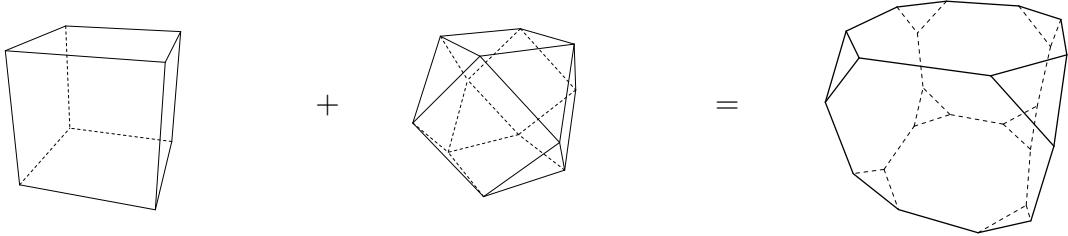
There is a straightforward implementation of f_{CELL} , which solves a sequence of LP's similar to (11.1) with objective functions y_1, y_2, y_3, \dots . This means we may have to solve $O(m)$ LP's in the worst case. Presently we don't know how to implement it in a more efficient manner.

Theorem 11.5 *There is an implementation of ReverseSearch($\text{Adj}_{CELL}, \delta_{CELL}, S_{CELL}, f_{CELL}$) for the cell enumeration problem with time complexity $O(m d \text{ LP}(d, m)|V_{CELL}|)$ and space complexity $O(m d)$.*

Proof. To prove this, first we recall that Theorem 9.2 says, the time complexity of ReverseSearch is $O(\delta t(\text{Adj})|V| + t(f)|E|)$. As we remarked earlier, $\delta_{CELL} = m$, $t(\text{Adj}_{CELL}) = O(\text{LP}(d, m))$, and $t(f_{CELL}) = O(m \text{ LP}(d, m))$. Since $|E_{CELL}| \leq d |V_{CELL}|$ holds for any arrangement (see [27]), the claimed time complexity follows. The space complexity is clearly same as the input size $O(m d)$. \blacksquare

12 Minkowski Additions of Polytopes

A zonotope is a very special Minkowski sum of polytopes, namely, a Minkowski of line segments. In this section, we study the complexity of Minkowski sums of polytopes P_1, \dots, P_k in \mathbb{R}^d and some algorithms for computing Minkowski sums of polytopes.



There are three basic variations of the problem. When input is H-polytopes and output is also H-polytope, Tiwary [50] has recently proved that the associated decision problem is NP-hard for $k = 2$. Here the *associated decision problem* is to test whether a given H-polytope P is the Minkowski sum of given H-polytopes P_1, \dots, P_k . When input is V-polytopes and output is H-polytope, the problem contains the representation conversion for polytopes as a special case ($k = 1$) whose complexity is still unknown. The last case when both input and output are V-polytopes is the only case for which an output-polynomial algorithm is known.

In this section, we present a compact output-polynomial algorithm for the last case. The algorithm is a natural extension of (the dual form of) the reverse search algorithm given in Section 11.1.

Faces, Minkowski Decomposition and Adjacency

For a polytope P and for any vector $c \in \mathbb{R}^d$, the set of maximizers x of the inner product $c^T x$ over P is denoted by $S(P; c)$. Thus each nonempty face of P is $S(P; c)$ for some c . We denote by $F(P)$ the set of faces of P , by $F_i(P)$ the set of i -dimensional faces, and by $f_i(P)$ the number of i -dimensional faces, for $i = -1, 0, \dots, d$. For each nonempty face F , the relatively open polyhedral cone of outer normals of P at F is denoted by $N(F; P)$. Thus, $c \in N(F; P)$ if and only if $F = S(P; c)$. The *normal fan* $N(P)$ of P is the cell complex $\{N(F; P) | F \in F(P)\}$ whose body is \mathbb{R}^d . If F is i -dimensional ($i = 0, 1, \dots, d$), the normal cone $N(F; P)$ is $(d-i)$ -dimensional. Thus the extreme points of P are in one-to-one correspondence with the full dimensional faces (which we call the *regions* or *cells*) of the complex.

Proposition 12.1 *Let P_1, P_2, \dots, P_k be polytopes in \mathbb{R}^d and let $P = P_1 + P_2 + \dots + P_k$. Then a nonempty subset F of P is a face of P if and only if $F = F_1 + F_2 + \dots + F_k$ for some face F_i of P_i such that there exists $c \in \mathbb{R}^d$ (not depending on i) with $F_i = S(P_i; c)$ for all i . Furthermore, the decomposition $F = F_1 + F_2 + \dots + F_k$ of any nonempty face F is unique.*

Proof. The equivalence follows directly from the obvious relation [30, Lemma 2.1.4]

$$S(P_1 + P_2 + \dots + P_k; c) = S(P_1; c) + S(P_2; c) + \dots + S(P_k; c) \text{ for any } c \in \mathbb{R}^d.$$

For the uniqueness, let F be a nonempty face with $F = S(P; c)$ for some c and let $F = F_1 + F_2 + \cdots + F_k$ be any decomposition. First, note that $F_i \subseteq S(P; c)$ for all i , because the value $c^T x$ for any $x \in F$ is the sum of the maximum values $c^T x_i$ subject to $x_i \in P_i$ for $i = 1, \dots, k$, and thus if $x \in F$ and $x = x_1 + x_2 + \cdots + x_k$ for $x_i \in F_i$, then $x_i \in S(P_i; c)$. Now suppose there exists F_i properly contained in $S(P_i; c)$. Let v be an extreme point of $S(P_i; c)$ not in F_i . Then there is a linear function $w^T x$ such that $w^T v$ is strictly greater than any value attained by $x \in F_i$. Now let x^* be any point attaining the maximum of $w^T x$ over the polytope $F_1 + F_2 + \cdots + F_{i-1} + F_{i+1} + \cdots + F_k$. Clearly $x^* + v \in F$ but this point cannot be in $F_1 + F_2 + \cdots + F_k$, a contradiction. This proves the uniqueness. \blacksquare

We refer the unique decomposition $F = F_1 + F_2 + \cdots + F_k$ of a nonempty face F as the *Minkowski decomposition*. Here, the dimension of F is at least as large as the dimension of each F_i . Thus we have the following.

Corollary 12.2 *Let P_1, P_2, \dots, P_k be polytopes in \mathbb{R}^d and let $P = P_1 + P_2 + \cdots + P_k$. A vector $v \in P$ is an extreme point of P if and only if $v = v_1 + v_2 + \cdots + v_k$ for some extreme point v_i of P_i and there exists $c \in \mathbb{R}^d$ with $\{v_i\} = S(P_i; c)$ for all i .*

For our algorithm to be presented in the next section, it is important to characterize the adjacency of extreme points in P .

Corollary 12.3 *Let P_1, P_2, \dots, P_k be polytopes in \mathbb{R}^d and let $P = P_1 + P_2 + \cdots + P_k$. A subset E of P is an edge of P if and only if $E = E_1 + E_2 + \cdots + E_k$ for some face E_i of P_i such that $\dim(E_i) = 0$ or 1 for each i and all faces E_i of dimension 1 are parallel, and there exists $c \in \mathbb{R}^d$ with $E_i = S(P_i; c)$ for all i .*

The following variation of the above is useful for the algorithm to be presented. The essential meaning is that the adjacency of extreme points is inherited from those of Minkowski summands.

Proposition 12.4 *Let P_1, P_2, \dots, P_k be polytopes in \mathbb{R}^d and let $P = P_1 + P_2 + \cdots + P_k$. Let u and v be adjacent extreme points of P with the Minkowski decompositions: $u = u_1 + u_2 + \cdots + u_k$ and $v = v_1 + v_2 + \cdots + v_k$. Then u_i and v_i are either equal or adjacent in P_i for each i .*

Proof. Let u and v be adjacent extreme points. It is sufficient to show that $[u, v] = [u_1, v_1] + [u_2, v_2] + \cdots + [u_k, v_k]$ and each $[u_i, v_i]$ is a face of P_i . Let $c \in \mathbb{R}^d$ be such that $[u, v] = S(P; c)$. Because $[u, v] = S(P_1; c) + S(P_2; c) + \cdots + S(P_k; c)$ and by the uniqueness of decomposition of u and v , both u_j and v_j are in $S(P_j; c)$, for all j . This implies that $[u_j, v_j] \subseteq S(P_j; c)$, for all j . On the other hand, one can easily see that in general $[u, v] \subseteq [u_1, v_1] + [u_2, v_2] + \cdots + [u_k, v_k]$. The last two relations give $[u_j, v_j] = S(P_j; c)$ for all j . This completes the proof. \blacksquare

This proposition immediately provides a polynomial algorithm for listing all neighbors of a given extreme point using linear programming.

12.1 Complexity of Minkowski Sums of V-Polytopes

The nontriviality of computing Minkowski sums of V-polytopes can be understood by how the complexity of Minkowski sums varies from some instances to another. In particular, we are most concerned with the complexity of sums in terms of the size of summands.

The first proposition shows that the vertex complexity of Minkowski sums is linearly bounded by the vertex complexity of summand polytopes.

Proposition 12.5 (Linearly Bounded Minkowski-Addition) . *For each $k \geq 2$ and $d \geq 2$, there is an infinite family of Minkowski additions for which $f_0(P_1 + P_2 + \cdots + P_k) \leq f_0(P_1) + f_0(P_2) + \cdots + f_0(P_k)$.*

Proof. Suppose $k \geq 2$ and $d \geq 2$. First pick up any d -polytope, say Q , with at least k extreme points, and select k extreme points. For each j th selected extreme point v^j , make a new polytope P_j from Q by truncating only v^j with one or more hyperplanes. Now we claim that the number $f_0(P_1 + P_2 + \cdots + P_k) \leq f_0(P_1) + f_0(P_2) + \cdots + f_0(P_k)$. See Figure 12.1 for an example for $k = 2$, $d = 3$ and Q is a 3-cube. To see this, let v be an extreme point of P_j for some fixed j . There are three cases. The first case is when v is an unselected one, i.e. an extreme point of Q not selected. In this case, it can be an Minkowski summand of an extreme point of P in a unique way, since any linear function maximized exactly at v over P_j is maximized exactly at v over other P_i 's. The second case is when v is a newly created vertex by the truncation of v^j . Since it is obtained by the truncation of v^j , any linear function maximized exactly at v over P_j is maximized exactly at v^j over other other P_i 's. The last case is when $v = v^i$ for some $i \neq j$. This case is essentially the same as the second case where v contributes uniquely to a new extreme point with each truncation vertex of P_i . By Corollary 12.2, every extreme point of P_j contributes at most once to $f_0(P_1 + P_2 + \cdots + P_k)$. This completes the proof. \blacksquare

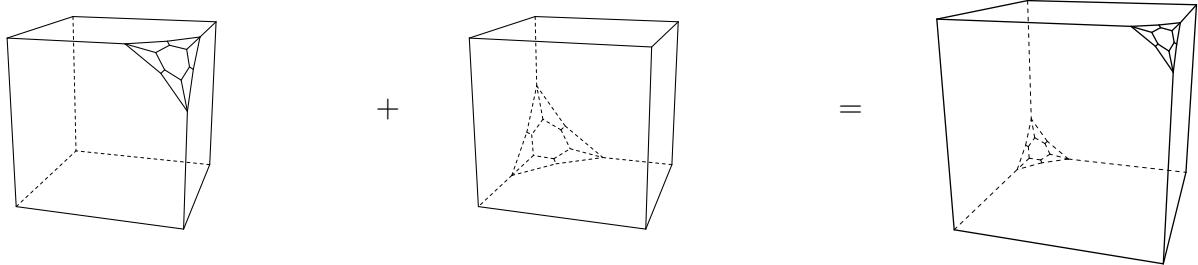


Figure 12.1: Minkowski Sum of Truncated Cubes

The following theorem gives the other extreme to the previous proposition. Namely, the obvious upper bound of the vertex complexity can be achieved for a large class of Minkowski sums of polytopes.

Theorem 12.6 (Tight Upper Bound [28]) .

In dimension $d \geq 3$, it is possible to choose k ($\leq d - 1$) polytopes P_1, \dots, P_k so that the trivial upper bound for the number of vertices is attained by their Minkowski sum.

$$f_0(P_1 + P_2 + \dots + P_k) = f_0(P_1) \times f_0(P_2) \times \dots \times f_0(P_k).$$

Proof. Here we give outline only, see [28] for a precise construction. On k ($\leq d - 1$) orthogonal planes in \mathbb{R}^d , place v_i points in convex position. Perturb the points slightly to make each P_i full dimensional. Figure 12.2 shows the case when $f_0(P_1) = f_0(P_2) = 4$ and $f_0(P) = f_0(P_1) \times f_0(P_2) = 16$ ■

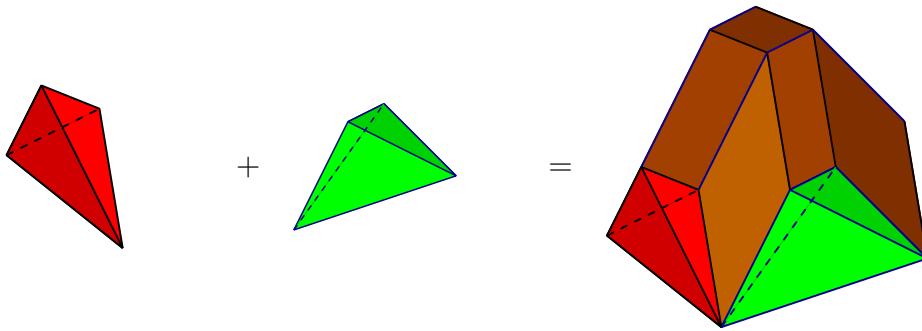


Figure 12.2: Minkowski Sum of $(d - 1)$ Thin Polytopes in Orthogonal Spaces

12.2 Extension of a Zonotope Construction Algorithm

We assume in this section that P_1, P_2, \dots, P_k are polytopes in \mathbb{R}^d given by the sets V_1, V_2, \dots, V_k of extreme points. We also assume that the graph $G(P_j)$ of P_j is given by the adjacency list $(\text{Adj}_j(v, i) : i = 1, \dots, \delta_j)$ of vertices adjacent to vertex $v \in V_j$ in graph $G(P_j)$, where δ_j is the maximum degree of $G(P_j)$ for each $j = 1, \dots, k$. If the degree $\deg_j(v)$ of v is less than δ_j in $G(P_j)$, we assume that $\text{Adj}_j(v, i) = \text{null}$ for all $i > \deg_j(v)$. Finally we define $\delta = \delta_1 + \delta_2 + \dots + \delta_k$, an upper bound of the maximum degree of $G(P)$, due to Proposition 12.4. For example, when the input polytopes are simple and full dimensional then $\delta_j = d$ for all j and $\delta = k d$. Note that for a given set V_j , one can compute the adjacency list in polynomial time using linear programming.

Recall that the Minkowski addition problem is to compute the set V of extreme points of $P = P_1 + P_2 + \dots + P_k$. We shall present a compact polynomial algorithm for the Minkowski addition problem.

The key idea in our algorithm design

The main algorithmic idea is quite simple. Just like for the vertex enumeration for convex polyhedra using reverse search given in Section 9.2, it traces a directed spanning tree T of the graph $G(P)$ of P rooted at an initial extreme point v^* . The difference from the vertex enumeration algorithm is that the polytope P is not given by a system of inequalities (i.e. not an H-polytope) in the present setting but as a Minkowski-addition of V-polytopes. Thus

we need to introduce a new way of defining a directed spanning tree that is easy to trace. We shall use the following simple geometric property of normal fans.

Proposition 12.7 *Let v and v' be two distinct extreme points of P , and let $c \in N(v; P)$ and $c' \in N(v'; P)$. Then there exists an extreme point v'' adjacent to v such that $N(v''; P)$ contains a point of form $(1 - \theta)c + \theta c'$ for some $0 \leq \theta \leq 1$.*

Proof. Since $v \neq v'$, their outer normal cones are two distinct full dimensional cones in the normal fan $N(P)$. This means that the parameterized point $t(\theta) := c + \theta(c' - c)$ ($0 \leq \theta \leq 1$) in the line segment $[c, c']$ must leave at least one of the bounding halfspaces of the first cone $N(v; P)$ as θ increases from 0 to 1. Since the bounding halfspaces of $N(v; P)$ are in one-to-one correspondence with the edges of G incident to v , any one of the halfspaces violated first corresponds to a vertex v'' adjacent to v claimed by the proposition. \blacksquare

Let us fix v^* as an initial extreme point of P . Finding one extreme point of P is easy. Just select any generic $c \in \mathbb{R}^d$, and find the unique maximizer extreme point v^i of $c^T x$ over P_i , for each i . The point $v = v^1 + v^2 + \cdots + v^k$ is an extreme point of P .

Now we construct a directed spanning tree of $G(P)$ rooted at v^* as follows. Let $v \in V$ be any vertex different from v^* . We assume for the moment that there is some canonical way to select an interior point of the normal cone of P at any given vertex, as we shall give one method to determine such a point later. Let c and c^* be the canonical vector of $N(v; P)$ and $N(v^*; P)$, respectively. By Proposition 12.7, by setting $v' = v^*$, we know that there is a vertex v'' adjacent to v such that $N(v''; P)$ meets the segment $[c, c^*]$. In general there might be several such vertices v'' (degeneracy). We break ties by the standard symbolic perturbation of c as $c + (\epsilon^1, \epsilon^2, \dots, \epsilon^d)^T$ for sufficiently small $\epsilon > 0$. Define the mapping $f : V \setminus \{v^*\} \rightarrow V$ as $f(v) = v''$. This mapping, called a *local search function* in reverse search, determines the directed spanning tree $T(f) = (V, E(f))$ rooted at v^* , where $E(f)$ is the set of directed edges $\{(v, f(v)) | v \in V \setminus \{v^*\}\}$.

Proposition 12.8 *The digraph $T(f)$ is a spanning tree of $G(P)$ (as undirected graph) and v^* is a unique sink node of $T(f)$.*

Proof. By the construction, v^* is a unique sink node of $T(f)$. It is sufficient to show that $T(f)$ has no directed cycle. For this, take any edge $(v, v'' = f(v)) \in E(f)$. Let c, c^* be the canonical vector for v, v^* , respectively. Without loss of generality, we assume nondegeneracy, since one can replace c with the perturbed vector $c + (\epsilon^1, \epsilon^2, \dots, \epsilon^d)^T$. Since c is an interior point of $N(v; P)$,

$$c^T(v - v'') > 0. \quad (12.1)$$

Again, by the construction and because the canonical points are selected as interior points of the associated normal cones, there exists $0 < \theta < 1$ such that $\hat{c} := (1 - \theta)c + \theta c^* \in N(v''; P)$. This implies $\hat{c}^T(v'' - v) > 0$, that is,

$$\begin{aligned} 0 &< ((1 - \theta)c + \theta c^*)^T(v'' - v) \\ &= (1 - \theta)c^T(v'' - v) + \theta(c^*)^T(v'' - v) \\ &< \theta(c^*)^T(v'' - v) \quad (\text{by (12.1)}). \end{aligned}$$

This implies that the vertex v'' attains a strictly higher inner product with c^* than v . Therefore, there is no directed cycle in $T(f)$. \blacksquare

Figure 12.3 shows an example of the directed spanning tree $T(f)$ in green.

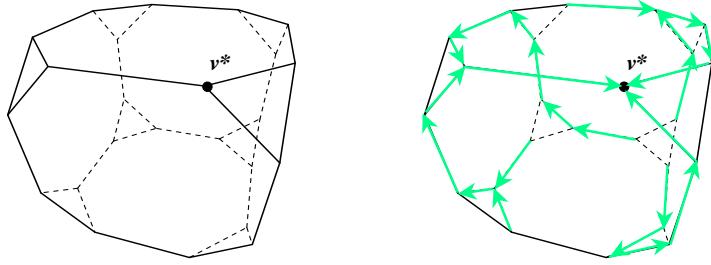


Figure 12.3: The Graph $G(P)$ and A Rooted Spanning Tree $T(f)$

A reverse search algorithm, to be presented below, traces reversely the tree from the root v^* in depth-first manner, using an adjacency oracle.

The critical computation in our algorithm is solving a linear programming problem. We denote by $\text{LP}(d, m)$ the time, as we used in Section 8. necessary to solve a linear programming in d variables and m inequalities.

Now we can state the complexity of our algorithm.

Theorem 12.9 *There is a compact polynomial algorithm for the Minkowski addition of k polytopes that runs in time $O(\delta \text{LP}(d, \delta) f_0(P))$ and space linear in the input size.*

The algorithm

The sequel of the section is devoted to present the technical details of a reverse search algorithm that traces $T(f)$ starting from its root vertex v^* against the orientation of edges. We shall prove Theorem 12.9 at the end.

As usual, our reverse search algorithm requires, in addition to the local search function f , an adjacency oracle function that implicitly determines the graph $G(P)$.

Let v be any vertex of P with the Minkowski decomposition $v = v_1 + v_2 + \dots + v_k$ (see, Corollary ??). Let

$$\Delta = \{(j, i) : j = 1, \dots, k \text{ and } i = 1, \dots, \delta_j\}. \quad (12.2)$$

Recall that for any $(j, i) \in \Delta$, $\text{Adj}_j(v_j, i)$ is the i th vertex adjacent to v_j whenever it is not *null*. We shall call a pair (j, i) *valid* for v if $\text{Adj}_j(v_j, i) \neq \text{null}$, and *invalid* otherwise. Let us define the associated edge vectors $e_j(v_j, i)$ by

$$e_j(v_j, i) = \begin{cases} \text{Adj}_j(v_j, i) - v_j & (j, i) \text{ is valid for } v \\ \text{null} & \text{otherwise.} \end{cases} \quad (12.3)$$

Proposition 12.4 shows that all edges of P incident to v are coming from the edges incident to v_j 's, or more precisely, each edge of P incident to v is parallel to some $e_j(v_j, i)$. This immediately implies that δ is an obvious upper bound of the degree of v . For each $(s, r) \in \Delta$, let us group the same (parallel) directions together as

$$\Delta(v, s, r) = \{(j, i) \in \Delta : e_j(v_j, i) \parallel e_s(v_s, r)\}. \quad (12.4)$$

Consider it as the empty set if (s, r) is invalid. Now, for any given pair $(s, r) \in \Delta$, checking whether $e_s(v_s, r)$ determines an edge direction of P is easily reducible to an LP (or more precisely, a linear feasibility problem):

$$\begin{aligned} e_s(v_s, r)^T \lambda &< 0, \\ e_j(v_j, i)^T \lambda &\geq 0 \quad \text{for all valid } (j, i) \notin \Delta(v, s, r). \end{aligned} \quad (12.5)$$

More precisely, the system (12.5) has a solution λ if and only if the direction $e_s(v_s, r)$ determines an edge of P incident to v . If it has a feasible solution, then by Proposition 12.4, the vertex \hat{v} adjacent to v along this direction is given by

$$\begin{aligned} \hat{v} &= \hat{v}_1 + \hat{v}_2 + \cdots + \hat{v}_k \\ \hat{v}_j &= \begin{cases} \text{Adj}_j(v_j, i) & \text{if there exists } i \text{ such that } (j, i) \in \Delta(v, s, r) \\ v_j & \text{otherwise.} \end{cases} \end{aligned}$$

Let us denote by $\Delta(v)$ as the set of all pairs $(s, r) \in \Delta$ such that $e_s(v_s, r)$ determines an edge of P and (s, r) is a member of $\Delta(v, s, r)$ with the smallest first index. This set represents a duplication-free index set of all edge directions at v .

Now we are ready to define our adjacency oracle as a function $\text{Adj} : V \times \Delta \rightarrow V \cup \{\text{null}\}$ such that

$$\text{Adj}(v, (s, r)) = \begin{cases} \hat{v} & \text{if } (s, r) \in \Delta(v) \\ \text{null} & \text{otherwise.} \end{cases} \quad (12.6)$$

Lemma 12.10 *One can evaluate the adjacency oracle $\text{Adj}(v, (s, r))$ in time $\text{LP}(d, \delta)$.*

Proof. The essential part of the evaluation is solving the system (12.5). Since $\delta = |\Delta|$, the system has d variables and at most δ inequalities and the claim follows. \blacksquare

Lemma 12.11 *There is an implementation of the local search function $f(v)$ with evaluation time $O(\text{LP}(d, \delta))$, for each $v \in V \setminus \{v^*\}$ with the Minkowski decomposition $v = v_1 + v_2 + \cdots + v_k$.*

Proof. The implementation of f essentially depends on how we define the canonical vector of the normal cone $N(v; P)$. Like in the adjacency oracle implementation, we use an LP formulation. Since the set of directions $e_j(v, i)$ for valid $(j, i) \in \Delta$ include all edge directions at v , the normal cone $N(v; P)$ is the set of solutions λ to the system

$$e_j(v_j, i)^T \lambda \leq 0 \quad \text{for all valid } (j, i) \in \Delta.$$

Since we need an interior point of the cone, we formulate the following LP:

$$\begin{aligned} \max & \quad \lambda_0 \\ \text{subject to} & \\ e_j(v_j, i)^T \lambda + \lambda_0 &\leq 0 \quad \text{for all valid } (j, i) \in \Delta \\ \lambda_0 &\leq K. \end{aligned} \quad (12.7)$$

Here K is any positive constant. Since v is a vertex of P , this LP has an optimal solution. We still need to define a unique optimal solution. For this, we use a very pragmatic definition: fix one deterministic algorithm and define the canonical vector as the unique solution returned by the algorithm. Since the number of variables is $d + 1$ and the number of inequalities is at most $\delta + 1$, the assumptions on LP implies the time complexity $O(\text{LP}(d, \delta))$ to compute the canonical vector. Note that for practical purposes, we should probably add bounding inequalities for λ to the LP (12.7) such as $-1 \leq \lambda_i \leq 1$ for all i to make sure that the optimal solution stays in a reasonable range. This does not change the complexity.

An execution of f requires to compute the canonical vectors c and c^* . Once they are computed, the remaining part is to determine the first bounding hyperplane of the normal cone $N(v; P)$ hit by the oriented line $t(\theta) := c + \theta(c^* - c)$ (as θ increases from 0 to 1). This amounts to solving at most δ one-variable equations, and is dominated by the canonical vector computation. \blacksquare

In Figure 12.4, we present the resulting reverse search algorithm, where we assume that the δ index pairs (j, i) in Δ are ordered as $(1, 1) < (1, 2) < \dots < (1, \delta_1) < (2, 1) < \dots < (k, \delta_k)$.

```

procedure MinkowskiAddition( $Adj, (\delta_1, \dots, \delta_k), v^*, f$ );
   $v := v^*$ ;  $(j, i) := (1, 0)$ ; (*  $(j, i)$ : neighbor counter *)
  output  $v$ ;
  repeat
    while  $(j, i) < (k, \delta_k)$  do
      increment  $(j, i)$  by one;
       $next := Adj(v, (j, i))$ ;
      if  $next \neq \text{null}$  then
        if  $f(next) = v$  then (* reverse traverse *)
           $v := next$ ;  $(j, i) := (1, 0)$ ;
          output  $v$ 
        endif
      endif
    endwhile;
    if  $v \neq v^*$  then (* forward traverse *)
       $u := v$ ;  $v := f(v)$ ;
      restore  $(j, i)$  such that  $Adj(v, (j, i)) = u$ 
    endif
  until  $v = v^*$  and  $(j, i) = (k, \delta_k)$ .

```

Figure 12.4: Reverse Search Algorithm for Minkowski Sum

Finally, we are ready to prove the main theorem, Theorem 12.9.

Proof. We use the general complexity result, Corollary 9.3, saying the time complexity of the reverse search in Figure 12.4 is $O(\delta(t(\text{Adj}) + t(f))|V|)$. By Lemma 12.10 and Lemma 12.11, both $t(\text{Adj})$ and $t(f)$ can be replaced by $\text{LP}(d, \delta)$. Since $f_0(P) = |V|$, the claimed

time complexity follows. The space complexity is dominated by those of the functions f and Adj which are clearly linear in the input size. \blacksquare

13 Problem Reductions in Polyhedral Computation

In this section, we look at some basic problems in polyhedral computation. Just like in combinatorial optimization, it is quite hard to distinguish hard problems (typically NP-hard problems) from easy problems. Here there are two sorts of easy problems. The first group consists of decision problems that are polynomially solvable. The second group consists of enumeration problems that may require output whose size is exponential in the input size, but are output-polynomially solvable.

We shall present some hard decision problems in Section 13.1, and discuss some hard enumeration problems in Section 13.2.

13.1 Hard Decision Problems in Polyhedral Computation

We start with two decision problems in polyhedral computation that are related to linear programming but are known to be hard.

For $A \in \mathbb{Q}^{m \times d}$ and $b \in \mathbb{Q}^m$, let $P_H(A, b)$ be the H-polyhedron

$$P_H(A, b) := \{x \in \mathbb{R}^d : Ax \leq b\}, \quad (13.1)$$

and let $\mathcal{A}(A, b)$ be the associated arrangement of hyperplanes:

$$\mathcal{A}(A, b) := \{h_1, \dots, h_m\}, \quad (13.2)$$

$$h_i := \{x \in \mathbb{R}^d : A_i x \leq b_i\}. \quad (13.3)$$

Problem 13.1 Optimal Vertex of a Polyhedron (OVP)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$, $c \in \mathbb{Q}^d$ and $K \in \mathbb{Q}$.
- Question: Does there exist a vertex v of $P_H(A, b)$ with $c^T v \geq K$?

Problem 13.2 K-Vertex of a Polyhedron (KVP)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$, $c \in \mathbb{Q}^d$ and $K \in \mathbb{Q}$.
- Question: Does there exist a vertex v of $P_H(A, b)$ with $c^T v = K$?

Theorem 13.3 ([26]) *The decision problems OVP and KVP are both NP-complete.*

Proof. It is clear that both problems are in the class NP. The proofs of the NP-completeness will be obtained by a polynomial time transformation from the following problem, known to be NP-complete in the strong sense [29]:

Problem 13.4 Directed Hamiltonian Path (DHP)

- Input: A directed graph $G = (V, A)$ and two distinct vertices $s, t \in V$.
- Question: Does G contain a directed Hamiltonian path from s to t ?

Let $G = (V, A)$ be a directed graph and $s \neq t \in V$. Associate a variable x_{ij} with each arc $(i, j) \in A$. Let $P(G)$ be the polytope given by:

$$\sum_{j|(i,j) \in A} x_{ij} - \sum_{j|(j,i) \in A} x_{ji} = 0, \quad \text{for each } i \in V - \{s, t\}, \quad (13.4)$$

$$\sum_{j|(s,j) \in A} x_{sj} - \sum_{j|(j,s) \in A} x_{js} = 1, \quad (13.5)$$

$$\sum_{j|(t,j) \in A} x_{tj} - \sum_{j|(j,t) \in A} x_{jt} = -1, \quad (13.6)$$

$$x_{ij} \geq 0, \quad \text{for each } (i, j) \in A. \quad (13.7)$$

The matrix of the coefficients of these inequalities is totally unimodular ([43], Proposition 2.6, p. 542) implying that $P(G)$ is integral. It follows that an extreme point x of $P(G)$ is the characteristic vector of a directed path joining s to t in G and, possibly, a set of circuits. If a circuit C exists, then x is a convex combination of the two points obtained by adding or subtracting small $\epsilon > 0$ on all the arcs of the circuit, a contradiction. Hence x is the characteristic vector of a simple directed path joining s to t . One verify easily that all such paths are extreme points of $P(G)$, proving that the extreme points of $P(G)$ are exactly the characteristic vectors of the simple directed paths joining s to t in G . These two facts thus imply that, for $K = |V| - 1$ and $c = \mathbf{1}$ (the vector of all 1's), both the OVP and the KVP problems for $P(G)$ are NP-complete in the strong sense. This completes the proof. \blacksquare

There are similar complexity results for arrangements of hyperplanes.

Problem 13.5 Optimal Vertex of an Arrangement (OVA)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$, $c \in \mathbb{Q}^d$ and $K \in \mathbb{Q}$.
- Question: Does there exists a vertex v of $\mathcal{A}(A, b)$ with $c^T v \geq K$?

Problem 13.6 Optimal Vertex of an Arrangement (KVA)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$, $c \in \mathbb{Q}^d$ and $K \in \mathbb{Q}$.
- Question: Does there exists a vertex v of $\mathcal{A}(A, b)$ with $c^T v = K$?

Theorem 13.7 ([26]) *The decision problems OVA and KVA are both NP-complete.*

Proof. Consider an instance of DHP and build a corresponding instance for OVA and KVA as follows: To each arc $(i, j) \in A$, we associate a variable x_{ij} . Let $d := |A|$, $K := |V| - 1$, $c = \mathbf{1}$ and define the arrangement generated by the following set of hyperplanes:

$$H_i := \{\mathbf{x} \in \mathbb{R}^d \mid \sum_{j|(i,j) \in A} x_{ij} - \sum_{j|(j,i) \in A} x_{ji} = 0\}, \quad \text{for each } i \in V - \{s, t\}, \quad (13.8)$$

$$H_s := \{\mathbf{x} \in \mathbb{R}^d \mid \sum_{j|(s,j) \in A} x_{sj} - \sum_{j|(j,s) \in A} x_{js} = 1\}, \quad (13.9)$$

$$H_t := \{\mathbf{x} \in \mathbb{R}^d \mid \sum_{j|(t,j) \in A} x_{tj} - \sum_{j|(j,t) \in A} x_{jt} = -1\}, \quad (13.10)$$

$$H_{ij} := \{\mathbf{x} \in \mathbb{R}^d \mid x_{ij} = 0\}, \quad \text{for each } (i, j) \in A. \quad (13.11)$$

First we observe that if DHP has a “yes” answer, so does the corresponding instance of OVA and KVA, as the characteristic vector of any directed Hamiltonian path lies on the $|V|$ hyperplanes H_i for $i \in V$ as well as on $(|A| - (|V| - 1)) = |A| - |V| + 1$ of the hyperplanes H_{ij} for $i \neq j \in V$. Note that the $|V|$ hyperplanes H_i for $i \in V$ are not linearly independent, but any subset of $(|V| - 1)$ of them are. Hence there are $(|A| - |V| + 1) + (|V| - 1) = |A|$ linearly independent hyperplanes containing the characteristic vector of any directed Hamiltonian path joining s to t in G , implying that the latter is a vertex of the given arrangement.

Now suppose that KVA or OVA has a “yes” answer produced by a vertex v of the constructed instance. One can write the $|A|$ equations defining the hyperplanes of the instance as a system of the form $Ax = b$. It is well known that the matrix $[A, b]$ is totally unimodular (see [43] for example). Thus any vertex of the arrangement has only $+1, -1$, or 0 coordinates, as shown by Cramer’s rule for solving a linear system.

Let S be a set of n linearly independent hyperplanes of the given family whose intersection is v . As the $|V|$ hyperplanes in $\{H_i \mid i \in V\}$ are not linearly independent, the number of these hyperplanes which are in S is at most $(|V| - 1)$. Hence the number of non zero coordinates of v is at most $(|V| - 1)$. As $c = \mathbf{1}$ and $c^T v \geq K = (|V| - 1)$, we have that exactly $(|V| - 1)$ coordinates of v are $(+1)$, all the others being (0) ’s. Thus v is the characteristic vector of a set P of $(|V| - 1)$ arcs of A . This also implies that KVA has a “yes” answer if and only if OVA has a “yes” answer.

If P is a directed Hamiltonian path in G joining s to t , then we are done. Otherwise, P contains a directed path joining s to t in G and at least one directed cycle C . But consider $v' \in R^n$ defined by

$$v'_{ij} = \begin{cases} 0 & \text{if } (i, j) \in C, \\ v_{ij} & \text{otherwise,} \end{cases} \quad \text{for each } (i, j) \in A. \quad (13.12)$$

This complete the proof. ■

13.2 Hard Enumeration Problems in Polyhedral Computation

For matrices $V \in \mathbb{Q}^{s \times d}$ and $R \in \mathbb{Q}^{t \times d}$, the V-polyhedron with representation pair (V, R) is denoted by $P_V(V, R)$, i.e.,

$$P_V(V, R) := \{x : x = V\lambda + R\mu, \mathbf{1}^T \lambda = 1, \lambda \geq \mathbf{0}, \mu \geq \mathbf{0}\}. \quad (13.13)$$

The following decision problem is arguably the most important problem in polyhedral computation.

Problem 13.8 Polyhedral Verification Problem (PVP)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$, $V \in \mathbb{Q}^{s \times d}$ and $R \in \mathbb{Q}^{t \times d}$.
- Question: Is $P_H(A, b) \neq P_V(V, R)$?

It is not difficult to prove that if PVP is in P, then there is an output-polynomial algorithm for the polyhedral representation conversion problem discussed in Section 9, see Polyhedral Computation FAQ [23].

PVP is easily seen to be in NP, because if the polyhedra are not equal, there is a succinct certificate for it, a point x in one of the polyhedra which is not in the other. Unfortunately, the complexity of PVP is still open. The decision problem PVP was first posed by Lovasz, see [49], and has been extensively studied by many researchers.

One of the most exciting progresses is the NP-completeness of a closely related problem, due to Khachiyan et al. [35].

Problem 13.9 Vertex Enumeration for an Unbounded Polyhedron (VEU)

- Input: $A \in \mathbb{Q}^{m \times d}$, $b \in \mathbb{Q}^m$ and $V \in \mathbb{Q}^{s \times d}$.
- Question: Does the H-polyhedron $P_H(A, b)$ contain a vertex not in V ?

Theorem 13.10 ([26]) *The decision problems VEU is NP-complete.*

Proof. (Outline) It is easily seen to be in NP, because if the answer is yes, then there is at least one vertex not in V . The proof uses a reduction from the NP-complete problem:

Problem 13.11 Negative Circuit Enumeration (NCE)

- Input: A digraph $G = (V, E)$ with edge weight $w : E \rightarrow \mathbb{Q}$, and a family S of negative circuits of G .
- Question: Does G contain a negative circuit not in the family S ?

Here, a *negative circuit* is a directed circuit $C \subset E$ whose total weight $\sum_{e \in C} w(e)$ is negative. It is shown by Khachiyan et al. [35] that NCE is NP-complete from a reduction from SAT. ■

(This section is to be extended.)

14 Evolutions and Applications of Polyhedral Computation

Polyhedral Computation has been shaped and polished through actual demands from numerous mathematicians, scientists, engineers and even social scientists. In this section, we present the author's personal involvements in various external or internal projects which have driven the advancement of polyhedral computation and software developments.

1987 – 1992: The First Stage. The first generation of codes for polyhedral representation conversion were written first for mathematicians to understand certain combinatorial polyhedra, such as **cut polytopes**, cut cones, and **traveling salesman polytopes**. It is extremely difficult to determine the facet inequalities of these polyhedra because typical associated combinatorial optimization problems are NP-hard. However, by computing the H-representation from a V-representation for small instances, many new facet inequalities were discovered and used for finding a stronger LP relaxation of NP-hard optimization problems. The first version of my implementation of the double description algorithm described in Section 9.1 was released in January 1988 is called **PDD** where p stands for the programming language Pascal. It helped the early stage of research on cut polytopes by Michel Deza and Monique Laurent, see [17, 18].

1993 – 1996: The Critical Second Stage. Then, a more computationally demanding task was needed for research in **material science**. Two physicists G. Ceder and G.D. Garbulsky at MIT contacted both David Avis and myself in 1993, and asked for our computational help in enumerating all extreme points of a highly degenerate polytope in dimension 8 given by 729 inequalities. The vertices represent physically stable states of a ternary (3 elements) alloy. David had a C-implementation named **RS** (which was replaced by **LRS** later) of the reverse search algorithm given in Section 9.2 then, and I had a C-version named **CDD** of the earlier code **PDD**. Both David and myself devoted our effort to compute the vertices, and finally it took us about a month to compute the results. Our greatest excitement came when we verified that the final results computed by our implementations of two totally different algorithms returned exactly the same results. This successful computation lead to a paper by the four of us [12].

About the same time, then a doctoral student Francisco Valero of **neuromuscular systems** laboratory at Stanford University contacted me. The application Valero discovered then surprised me considerably. I did not imagine that one can find an application of polyhedral computation in human bodies and muscles. He email in May 1994 describing his application reads

My application deals with **finger muscles having force limits** (i.e., from zero force to their maximum physiological force for each finger muscle) which defines a hypercube in a dimension equal to the number of muscles under consideration. Other mechanical, functional or anatomical characteristics produce further constrain equations (i.e., need the force of the finger to be zero in certain directions, finger size/configuration, etc.). The vertex enumeration technique helps me identify the limits of muscle force production

under these constraints, which in turn maps into functional limits such as maximum finger forces, accelerations, etc, which are independent of muscle force coordination. Coordination can be studied with standard linear programming techniques. The limits of function, however, require the explicit enumeration of the vertices of convex polyhedra in n -dimensional space.

Valero has been a strong advocate of computational geometry techniques applied to **biomedical** and **biomechanics** fields since then. A recent paper [36] shows the **analysis of muscle redundancies** using the vertex enumeration in polyhedra.

From the software development front, a new C++ version of CDD, called CDD+, was released in April 1995 which has the capability of using both floating-point and rational exact arithmetic using GMP [1].

1997 – 2007: Developments of Polyhedral Computation Libraries. Further advancements were made during this period for the development of software C-libraries CDDLIB and LRSLIB, based on CDD and LRS, respectively by Fukuda [22] and Avis [2]. Naturally, these libraries have been integrated into other programs.

A versatile **R-interface** of CDDLIB was written by the statistician Charles Geyer of University of Minnesota. It is available from

<http://www.stat.umn.edu/~charlie/>

A webpage of computing **all Nash equilibria** of bimatrix games using LRSLIB written by Rahul Savani became available at

<http://banach.lse.ac.uk/form.html>.

A **Matlab toolbox** for the study of **control theory** with an emphasis on parametric optimization was written by a group of researchers at the system dynamics and control group at ETH Zurich. It has an interface called CDDMEX to cddlib and is available at <http://control.ee.ethz.ch/research/software.en.html>

A **Python** interface PYPOLYHEDRON to cddlib was written by Pearu Peterson. He wrote in his email in 2007 “I am using it for analyzing multi-soliton interactions. In terms of computational geometry, I just construct special polyhedron in $(N + 1)$ -D space, project it to N -D space, and then find its intersection with 2-D hyperplane, which after projecting to 2-D space gives an interaction pattern of the N -soliton solution.” It is available at <http://cens.ioc.ee/projects/polyhedron/>

Polymake is a platform to do polyhedral and algebraic computation mainly for mathematicians whose two core engines are CDDLIB and LRSLIB. It is available at <http://www.polymake.org/doku.php>

TOPCOM [45] is a package for computing Triangulations Of Point Configurations and Oriented Matroids. It uses the LP code of CDDLIB for the recognition of regular triangulations.

Minksum [<http://www.cs.dartmouth.edu/~weibel/minksum.php>] is a program to compute the V-representation (i.e. the set of vertices) of the Minkowski addition of several convex polytopes given by their V-representation. It is an implementation in C++

language of the reverse search algorithm given in Section 12.2 whose time complexity is polynomially bounded by the sizes of input and output. It relies on the exact LP solver of CDDLIB.

Gfan [33] is a program to list all reduced **Gröbner bases** of a general polynomial ideal given by a set of generating polynomials in n -variables. It is an implementation in C++ language of the reverse search algorithm [25]. It relies on the exact LP solver of CDDLIB.

2004 – 2011: Expanding Application Fields. An application of **Minkowski sum of polytopes** presented in Section 12 is given in a doctoral thesis of J.P. Petit [44] in 2004, which is **computer aided tolerancing in design and manufacturing** using a mathematical model with convex polytopes in dimension 6. The dimension is simply $3 + 3$ where the first 3 is the dimension of the space and the latter is the freedom of movement in 3-space.

A polyhedral model was introduced in a doctoral research at Queen's University Belfast guided by Cecil Armstrong on **aircraft stress load evaluation and optimization**. The essential problem is to detect the most critical parts of aircrafts against a set of many likely stresses, which is reduced to the redundancy removal in linear inequality systems, the theme of Section 8.2.

To analyze the **effects of United Nations peacekeeping operations**, the danger of using the high dimensional analysis is pointed out in a paper by political scientists in [46], after a few researchers in computational geometry including myself presented counter-intuitive facts in higher dimensional spaces. In particular, one serious problem of estimating the effect of a future operation, a relatively small set of past instances represented by high dimensional points cannot be a reliable guidance, due to the fact that a new point will most likely be (far) outside of the **convex hull of the past data points**, and thus a wild extrapolation occurs at a high probability.

Future. From my personal involvements in polyhedral computation during the past 24 years, once reliable and efficient codes of polyhedral computation become available, new users might show up from any field of science, engineering, humanities and even arts. Thus, the main purpose of writing this book is to present the fundamental theory of polyhedra and the basic computational problems associated with polyhedra with most efficient algorithmic techniques. My belief is that interesting applications should follow once researchers have easy access to the theory and computational codes. After all, convex polyhedra appear almost everywhere, implicitly or explicitly.