Combinatorial and discrete geometry

2025/26 autumn
(updated: October 6, 2025)

BMETE94MM02

Contents

1 Convexity

1.1 Affine space

We will regard $\mathbb{R}^n$ as a Euclidean space with the *inner product* $\langle x, y\rangle = x_1y_1 + \cdots + x_ny_n$, which induces the (Euclidean) *norm* $\|v\| = \sqrt{\langle v, v\rangle}$. The Euclidean affine space is the affine space over $\mathbb{R}^n$. If we choose a point o in the affine space, then the points can be identified by their position vectors: the point p corresponds to the unique vector v that satisfies $p = o + x$. In this case we refer to the chosen point as the *origin*.

The *distance* of a pair of points p and q is $\operatorname{dist}(p, q) = \|q - p\|$. The Euclidean affine space with dist is a metric space. The *interior*, *boundary*, *closure* and *cardinality* of a set $X \subseteq \mathbb{R}^n$ will be denoted by $\operatorname{int}(X), \operatorname{bd}(X), \operatorname{cl}(X), |X|$, respectively.

Definition 1.1. The *Minkowski sum* of two subsets $X_1, X_2 \subseteq \mathbb{R}^n$ is the set

$$\{x_1 + x_2 | x_1 \in X_1, x_2 \in X_2\} .$$

The *multiple* of a subset $X \subseteq \mathbb{R}^n$ by $\lambda \in \mathbb{R}$ is the set $\{\lambda x | x \in X\}$.

It should be noted that both the Minkowski sum and the multiple depend on the origin, since the operations are defined at the vector space level. However, choosing a different origin only affects the results by a translation. Moreover, certain combinations such as $\frac{1}{2}X_1 + \frac{1}{2}X_2$ do not depend on the origin.

Definition 1.2 (affine subspace, dimension)**.** An affine subspace of $\mathbb{R}^n$ is a subset of the form $p + L$, where $p \in \mathbb{R}^n$ and $L \subseteq \mathbb{R}^n$ is a linear subspace. The dimension of the affine subspace $p + L$ is $\dim L$. Affine subspaces of dimension 1, 2, $n - 1$ are called lines, planes, and hyperplanes, respectively.

Remark 1.3. *If $p, p' \in \mathbb{R}^n$ and $L, L' \leq \mathbb{R}^n$ are linear subspaces, then $p + L = p' + L'$ iff $p - p' \in L = L'$.*

Let $u \in \mathbb{R}^n$ be a nonzero vector and $t \in \mathbb{R}$, and consider the sets

$$\begin{aligned} H_+ &= \{v \in \mathbb{R}^n | \langle v, u\rangle > t\} \\ H_0 &= \{v \in \mathbb{R}^n | \langle v, u\rangle = t\} \\ H_- &= \{v \in \mathbb{R}^n | \langle v, u\rangle < t\} . \end{aligned}$$

Then H_0 is a hyperplane, H_+ and H_- are the two *open half spaces* determined by H_0, and $H_0 \cup H_+$ and $H_0 \cup H_-$ are the two *closed half spaces*. The boundary of each of these half spaces is H_0. u is a *normal vector* of H_0 and also an outer (inner) normal vector of H_- (H_+).

Definition 1.4. Let $G_1 = p_1 + L_1$ and $G_2 = p_2 + L_2$ be affine subspaces of $\mathbb{R}^n$. G_1 and G_2 are *perpendicular* (or *orthogonal*) if for all vectors $v_1 \in L_1$ and $v_2 \in L_2$ we have $\langle v_1, v_2\rangle = 0$. G_1 and G_2 are *parallel* if $L_1 = L_2$.

Proposition 1.5. *A nonempty intersection of affine subspaces is an affine subspace.*

Definition 1.6. Let $X \subseteq \mathbb{R}^n$. The *affine hull* of X is the intersection of all affine subspaces that contain X. The affine hull of X is denoted by $\operatorname{aff}(X)$. The linear hull is $\operatorname{lin}(X) = \operatorname{aff}(X \cup \{o\})$.

Considering X as a subspace of $\operatorname{aff}(X)$ (which in turn inherits its topology from $\mathbb{R}^n$), we define the *relative interior* and *relative boundary* of X as the interior and boundary in $\operatorname{aff}(X)$. We use the notations $\operatorname{relint}(X)$ and $\operatorname{relbd}(X)$.

Definition 1.7. A set $X \subseteq \mathbb{R}^n$ is *affinely independent* if for all $x \in X$ we have $\operatorname{aff}(X \setminus \{x\}) \neq \operatorname{aff}(X)$. If X is not affinely independent, then we say that it is *affinely dependent.*

Definition 1.8. Let $p_1, p_2, \ldots, p_k \in \mathbb{R}^n$ and $\lambda_1, \lambda_2, \ldots, \lambda_k \in \mathbb{R}$ such that $\sum_{i=1}^k \lambda_i = 1$. Then the point $\sum_{i=1}^k \lambda_i p_i$ is the *affine combination* of the points $p_1, p_2, \ldots, p_k$ with *coefficients* $\lambda_1, \lambda_2, \ldots, \lambda_k$.

Proposition 1.9. *The affine hull of X is the set of all affine combinations of finite subsets of X.*

Corollary 1.10. *X is affinely independent if and only if no point $x \in X$ can be written as an affine combination of some other points from X.*

Theorem 1.11. *Let $X = \{p_1, \ldots, p_k\} \subset \mathbb{R}^n$. Then X is affinely independent if and only if $\sum\limits_{i=1}^{k} \lambda_i p_i = 0$ and $\sum\limits_{i=1}^{k} \lambda_i = 0$ implies $\lambda_i = 0$ for all values of i.*

Corollary 1.12. *If $X \subset \mathbb{R}^n$ is affinely independent, then every point of* $\operatorname{aff}(X)$ *can be uniquely written as an affine combination of some points in X.*

Theorem 1.13. *If $X \subseteq \mathbb{R}^n$ and $|X| \geq n+2$, then X is affinely dependent.*

Corollary 1.14. *Every affine subspace of the space $\mathbb{R}^n$ is the affine hull of a most $n+1$ points.*

Example 1.15. *The standard unit vectors and o form an affinely independent set of $n+1$ points in $\mathbb{R}^n$.*

1.2 Convex sets

Definition 1.16. Let $p_1, p_2, \ldots, p_k \in \mathbb{R}^n$ and $\lambda_1, \lambda_2, \ldots, \lambda_k \in \mathbb{R}$ such that $\sum_{i=1}^k \lambda_i = 1$ and $\lambda_i \geq 0$ for all i. Then the point $\sum_{i=1}^k \lambda_i p_i$ is the *convex combination* of the points $p_1, p_2, \ldots, p_k$ with *coefficients* $\lambda_1, \lambda_2, \ldots, \lambda_k$.

The set of all convex combinations of the points $p, q \in \mathbb{R}^n$ is the *closed segment* with endpoints p and q, denoted by $[p,q]$. If $p \neq q$, then $(p,q) = [p,q] \setminus \{p,q\}$ is the *open segment* with endpoints p and q.

Definition 1.17. Let $K \subseteq \mathbb{R}^n$. The set K is called *convex*, if for arbitrary $p, q \in K$ we have $[p,q] \subseteq K$.

Remark 1.18. *The intersection of a family of convex sets is convex.*

Definition 1.19. Let $X \subseteq \mathbb{R}^n$. The *convex hull* of X is the intersection of all convex sets that contain X. The convex hull of X is denoted by $\operatorname{conv}(X)$.

Theorem 1.20. *Let $K \subseteq \mathbb{R}^n$ be a closed convex set. Then K is equal to the intersection of the closed half spaces that contain K.*

Remark 1.21. *The closure of a convex set is convex.*

Corollary 1.22. *If K is convex and $x \in \operatorname{bd}(K)$, then there exists a hyperplane H such that $x \in H$ and K is a subset of one of the closed half spaces bounded by H.*

Definition 1.23. Let $K \subseteq \mathbb{R}^n$ be a convex set. If H is a closed half space satisfying $K \subseteq H$ and whose boundary intersects the boundary of K, we say that H is a *supporting half space* of K, and the boundary of H is a *supporting hyperplane* of K.

Theorem 1.24. *The convex hull of X is the set of all convex combinations of finite subsets of X.*

Definition 1.25. The convex hulls of k-element subsets of $\mathbb{R}^n$ with $k \le n+1$ are called *simplices*. If the point set is affinely independent, we call the simplex *nondegenerate*. Then the elements of the point set are the *vertices* of the nondegenerate simplex, and the convex hull of two vertices is an *edge* of the simplex. If $k = n+1$, then the convex hull of n vertices is a *facet* of the simplex. If all edges of a nondegenerate simplex are of equal length, we call the simplex *regular*.

Example 1.26. *The convex hull of the standard unit vectors in $\mathbb{R}^{n+1}$ is an n-dimensional regular simplex.*

Proposition 1.27. *Let H be a closed half space bounded by the hyperplane H_0, and let $X \subset H$ be arbitrary. Then $\operatorname{conv}(X) \cap H_0 = \operatorname{conv}(X \cap H_0)$.*

We continue with the fundamental theorems of convex geometry.

Theorem 1.28 (Radon)**.** *Let $X \subset \mathbb{R}^n$ be a set containing at least $n+2$ points. Then X can be decomposed into two parts whose convex hulls have a nonempty intersection.*

Theorem 1.29 (Carathéodory)**.** *Let $X \subset \mathbb{R}^n$ be a nonempty set. If $p \in \operatorname{conv}(X)$, then X has a subset Y consisting of at most $n+1$ points, satisfying $p \in \operatorname{conv}(Y)$.*

Theorem 1.30 (colorful Carathéodory theorem)**.** *Let $X_1, X_2, \ldots, X_{n+1} \subset \mathbb{R}^n$ be compact subsets. Assume that for any i we have $o \in \operatorname{conv} X_i$. Then there exist points $p_i \in X_i$ such that o is contained in $\operatorname{conv}\{p_1, p_2, \ldots, p_{n+1}\}$.*

In the theorem, X_i denotes the set of points with 'color i'. Thus, the statement guarantees that there is a 'rainbow simplex' containing the origin.

Theorem 1.31 (Helly, finite)**.** *Let $\mathcal{K}$ be a finite family of at least $n+1$ convex sets in $\mathbb{R}^n$. If any $(n+1)$ elements of $\mathcal{K}$ have a nonempty intersection, then all elements of $\mathcal{K}$ have a nonempty intersection.*

Theorem 1.32 (Helly, infinite)**.** *Let $\mathcal{K}$ be a family of at least $n+1$ closed, convex sets in $\mathbb{R}^n$ such that at least one member of $\mathcal{K}$ is compact. Assume that any $n+1$ elements of $\mathcal{K}$ have a nonempty intersection. Then there is a point which is contained in every element of $\mathcal{K}$.*

The following statement can be proved using Carathéodory's theorem.

Theorem 1.33. *Let $H \subset \mathbb{R}^n$ be compact. Then* $\operatorname{conv}(H)$ *is also compact.*

Exercise 1.1. What are the possible intersections of two planes in $\mathbb{R}^4$?

Exercise 1.2. Let K be a convex set and L a line such that $L \cap \operatorname{int} K \neq \emptyset$. What are the possible cardinalities of $L \cap \operatorname{bd} K$?

Exercise 1.3. Show that the interior of a convex set is convex.

Exercise 1.4. Which of the following properties are preserved by conv?

(i) finite

(ii) bounded

(iii) closed

(iv) open

Exercise 1.5. Prove that for arbitrary subsets $X, Y \subseteq \mathbb{R}^n$ the equality $\operatorname{conv}(X+Y) = \operatorname{conv}(X) + \operatorname{conv}(Y)$ holds.

Exercise 1.6. Define the map $I : 2^{\mathbb{R}^n} \to 2^{\mathbb{R}^n}$ as $I(X) = \bigcup_{x,y \in X} [x, y]$. Prove that $I^{\lceil \log_2(n+1) \rceil}(X) = \operatorname{conv} X$.

Exercise 1.7. Can the $n+1$ be improved to n in Helly's theorem?

Exercise 1.8. Show by examples that the conditions in the infinite version of Helly's theorem are necessary: the implication may fail if the sets are not closed or there is no compact member.

Exercise 1.9. Given a Borel probability measure μ on $\mathbb{R}^n$, a point $x \in \mathbb{R}^n$ is a *centerpoint* if every closed half space that contains x has measure at least $\frac{1}{n+1}$. Prove that every Borel probability measure has at least one centerpoint.

2 Convex polytopes

2.1 Polytopes and polyhedral sets

Consider a convex polygon in the plane (Figure 1), and let its vertices in cyclic order be $x_1, \ldots, x_n$. Then the segments $[x_1, x_2], [x_2, x_3], \ldots, [x_{n-1}, x_n], [x_n, x_1]$ are its edges, and each edge determines a half plane that contains the polygon and whose boundary contains the edge. The polygon is the intersection of these half planes. In particular, while a general convex closed set in $\mathbb{R}^2$ is the intersection of closed half planes, a convex polygon is the intersection of *finitely many* closed half planes.

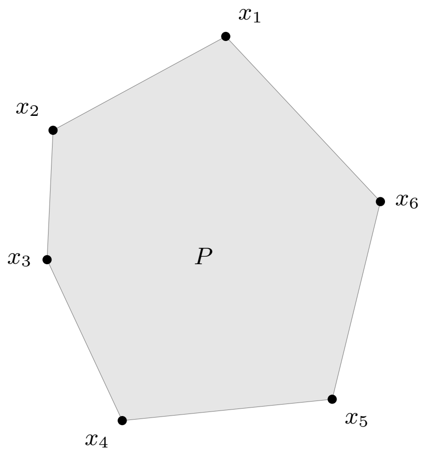

Figure 1: A convex polygon.

These two dual descriptions of convex polygons suggest two possible generalizations to higher dimensional spaces: convex hulls of finite sets of points and intersections of finitely many closed half spaces. These two classes of convex sets are not the same even in the plane, since the convex hull of finitely many points is bounded, while the intersection of finitely many half planes (e.g., of just one or even zero of them) can be unbounded. But this is the only difference: any *bounded* set that can be written as the intersection of finitely many closed half-planes is indeed a polygon.

In the following we study the generalizations of both descriptions to higher dimensions. We consider both the bounded and the unbounded versions. For the unbounded case we will use the concept of a cone.

Definition 2.1. A set of vectors $C \subseteq \mathbb{R}^d$ is a *cone* if $0 \in C$, $C + C \subseteq C$ and for all $\lambda \geq 0$ we have $\lambda \cdot C \subseteq C$.

In other words, any linear combination of vectors from C with nonnegative coefficients is also in C.

Definition 2.2. The *conical hull* of a set $Y \subseteq \mathbb{R}^d$ is the intersection of all cones in $\mathbb{R}^d$ that contain Y. The conical hull of Y is denoted by $\text{cone}(Y)$

The situation is similar to that of linear, affine, and convex hulls: the conical hull is a cone, and it can be shown that $\text{cone}(Y)$ is the set of linear combinations of vectors from Y with nonnegative coefficients.

Definition 2.3. An *H-polyhedral set* in $\mathbb{R}^d$ is the intersection of finitely many closed half-spaces. An *H-polytope* is a bounded H-polyhedral set.

A *V-polytope* in $\mathbb{R}^d$ is the convex hull of finitely many points. A *V-polyhedral set* is a set of the form $P = \operatorname{conv}(V) + \operatorname{cone}(Y)$ where $|V| < \infty$ and $|Y| < \infty$, i.e., the Minkowski sum of a V-polytope and the conical hull of a finite set of vectors.

By definition, an H-polyhedral set is bounded if and only if it is an H-polytope.

Proposition 2.4. *An V-polyhedral set is bounded if and only if it is a V-polytope.*

Proof. A V-polytope is bounded and every V-polytope is a V-polyhedral set (take $Y = \emptyset$ in the definition), therefore we only need to show that a bounded V-polyhedral set is a V-polytope.

Let $P = \operatorname{conv}(V) + \operatorname{cone}(Y)$ be bounded. If V is empty, then P is also empty, therefore it is a V-polytope. Otherwise $\operatorname{cone}(Y)$ must be bounded, since it is a subset of $P + (-1) \cdot \operatorname{conv}(V)$. But then $Y = \{0\}$, otherwise it would contain some $y \neq 0$ and λy for all $\lambda \geq 0$, an unbounded subset. □

Theorem 2.5. *Let $P \subseteq \mathbb{R}^d$ be a subset. Then*

(i) P is a V-polyhedral set if and only if it is an H-polyhedral set.

(ii) P is a V-polytope if and only if it is an H-polytope.

By the above theorem, from a mathematical point of view, V-polyhedral sets and H-polyhedral sets are the same classes of subsets of $\mathbb{R}^d$, and likewise V-polytopes and H-polytopes are the same. In alignment with the usual terminology, we will call such sets polytopes and polyhedral sets, respectively.

It is useful to have both kinds of characterizations, because statements about polytopes and polyhedral sets such as

(i) The intersection of a polytope and a polyhedral set is a polytope. (H)

(ii) The Minkowski sum of two polytopes is a polytope. (V)

(iii) Every projection of polytope is a polytope. (V)

can be much easier to prove using one or the other.

On the other hand, from a computational perspective, it is useful to distinguish between these concepts. If we think of a polyhedral set as the input or output of an algorithm, then it makes a difference in terms of complexity whether it is described as a convex hull of a finite set of points of the intersection of a finite set of half-spaces. One reason is that the complexity is usually measured as a function of the input size (which can be very different depending on the description), but also the "work" required to solve a problem can be very different.

As an example, consider maximizing a given linear functional over a polytope. If the polytope is specifed as a convex hull of a finite set of points, then the task is to evaluate the functional at these points and find the largest value obtained. On the other hand, maximizing a linear functional on the intersection of finitely many half spaces is the same as maximization subject to linear constraints, the topic of linear programming.

We list some of the most common polyhedral sets.

Example 2.6.

(i) Every simplex (Definition 1.25) is a polytope.

(ii) The d-dimensional cube $[-1,1]^d \subseteq \mathbb{R}^d$ *is a polytope.*

(iii) The cross polytope *is the convex hull of the points* $\pm e_1, \pm e_2, \dots, \pm e_d$*, where* $e_1, \dots, e_d$ *are the standard basis vectors. As the name suggests, the cross polytope is a polytope.*

(iv) The nonnegative orthant *is the set* $\{(x_1, x_2, \dots, x_d) | x_1 \geq 0, x_2 \geq 0, \dots, e_d \geq 0\}$*. The nonnegative orthant is a polyhedral set (and also a cone).*

Recall that a supporting hyperplane of a convex set K is a hyperplane H that bounds a half space H_+ such that $K \subseteq H_+$ and $K \cap H \neq \emptyset$.

Definition 2.7. Let $K \subseteq \mathbb{R}^d$ be a closed, convex set. A *face* of K is

(i) the intersection of K with a supporting hyperplane, or

(ii) K itself, or

(iii) the empty set.

The *dimension* of a face is the dimension of its affine hull. A face of dimension j is also called a j-face. Special names are in use for faces of certain dimensions: a 0-face is a *vertex*, a 1-face is an *edge*, and if $\dim K = d$, then a $(d-1)$-face is a *facet*.

Theorem 2.8. *Let $P \subseteq \mathbb{R}^d$ be a polytope, and let V be the set of its vertices. Then*

(i) $P = \mathrm{conv}(V)$*.*

(ii) If F is a face of P, then F is a polytope and the set of its vertices is $F \cap V$.

(iii) If F_1 and F_2 are faces of P, then $F_1 \cap F_2$ is a face of P as well.

(iv) The faces of F are the faces of P that are contained in F

Lemma 2.9. *Let $P \subset \mathbb{R}^d$ be a d-dimensional polytope and let $x \in \mathrm{bd}(P)$ be arbitrary. Then there is a unique face of P containing x in its relative interior.*

Exercise 2.1. A convex subset of $\mathbb{R}^n$ is called *locally polyhedral* (also known as *quasi-polyhedral* or *boundedly polyhedral*), if its intersection with every polytope is a polytope (possibly empty). Given an example of a locally polyhedral set in $\mathbb{R}^2$ which is not polyhedral.

Exercise 2.2. The standard d-simplex is the convex hull of the $d+1$ standard unit vectors in $\mathbb{R}^{d+1}$. Express it as the intersection of finitely many half spaces.

Exercise 2.3. Describe the cube $[-1,1]^d$ as a V-polytope and as an H-polytope.

Exercise 2.4. Describe the cross polytope as a V-polytope and as an H-polytope.

2.2 Euler characteristic

Definition 2.10. The indicator function of a subset $A \subseteq \mathbb{R}^d$ is the function $I[A]$ defined as

$$I[A](x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

It follows direcly from the definition that $I[A] \cdot I[B] = I[A \cap B]$ for all $A, B \subseteq \mathbb{R}^d$.

Lemma 2.11 (Inclusion-exclusion formula)**.** *For any sets $A_1, A_2, \dots, A_n \subset \mathbb{R}^d$,*

$$\begin{aligned} I[A_1 \cup A_2 \cup \ldots \cup A_n] &= 1 - (1 - I[A_1])(1 - I[A_2]) \ldots (1 - I[A_n]) \\ &= \sum_{j=1}^{n} (-1)^{j-1} \sum_{1 \le i_1 < i_2 < \ldots < i_j \le n} I[A_{i_1} \cap A_{i_2} \cap \ldots \cap A_{i_j}]. \end{aligned}$$

Proof. Introducing the notation $\overline{B} = \mathbb{R}^d \setminus B$ for any set $B \subseteq \mathbb{R}^d$, the first statement is equivalent to the equality

$$A_1 \cup A_2 \cup \ldots \cup A_n = \overline{\overline{A_1} \cap \overline{A_2} \cap \ldots \overline{A_n}},$$

which readily follows from the de Morgan identities. The second statement is a consequence of the previous remark. □

Definition 2.12. The real vector space generated by the indicator functions $I[A]$ of the compact, convex sets $A \subset \mathbb{R}^d$ is called the *algebra of compact, convex sets*, and is denoted by $\mathcal{K}(\mathbb{R}^d)$. The real vector space generated by the indicator functions $I[A]$ of the closed, convex sets $A \subset \mathbb{R}^d$ is called the *algebra of closed, convex sets*, and is denoted by $\mathcal{C}(\mathbb{R}^d)$.

Remark 2.13. *An arbitrary element of $\mathcal{K}(\mathbb{R}^d)$ can be written as $\sum_{i=1}^n \alpha_i I[A_i]$, where $\alpha_i \in \mathbb{R}$, and the sets $A_i \subset \mathbb{R}^d$ are compact and convex. Observe that if $A, B \subset \mathbb{R}^n$ are compact, convex sets, then $A \cap B$ is also compact and convex, implying that the product of two elements of $\mathcal{K}(\mathbb{R}^d)$ is also an element of $\mathcal{K}(\mathbb{R}^d)$. Thus, the set $\mathcal{K}(\mathbb{R}^d)$ is indeed an algebra over $\mathbb{R}$. A similar observation can be made about the algebra $\mathcal{C}(\mathbb{R}^d)$.*

Definition 2.14. A linear map $\mathcal{K}(\mathbb{R}^d) \to \mathbb{R}$ or $\mathcal{C}(\mathbb{R}^d) \to \mathbb{R}$ is called a *valuation.*

Theorem 2.15. *There is a unique valuation $\chi : \mathcal{C}(\mathbb{R}^d) \to \mathbb{R}$ satisfying $\chi(I[A]) = 1$ for all nonempty, closed, convex sets $A \subset \mathbb{R}^d$.*

This valuation is called the *Euler characteristic* induced on the algebra of closed, convex sets. Theorem 2.15 was first proved by H. Hadwiger.

The following lemma is a consequence of Lemma 2.11 and Theorem 2.15.

Lemma 2.16. *Let $A_1, A_2, \ldots, A_n \subset \mathbb{R}^d$ be sets such that $I[A_i] \in \mathcal{K}(\mathbb{R}^d)$ for any $i = 1, 2, \ldots, n$. Then*

$$\chi(A_1 \cup A_2 \cup \ldots \cup A_n) = \sum_{j=1}^{n} (-1)^{j-1} \sum_{1 \leq i_1 < i_2 < \ldots < i_j \leq n} \chi(A_{i_1} \cap A_{i_2} \cap \ldots \cap A_{i_j}).$$

Lemma 2.17. *Let $P \subset \mathbb{R}^d$ be a d-dimensional (convex) polytope. Then*

$$\chi(\operatorname{bd} P) = 1 + (-1)^{d-1}, \quad \text{and} \quad \chi(\operatorname{int} P) = (-1)^d.$$

Given a d-dimensional polytope P, let $f_j(P)$ denote the number of j-faces of P.

Theorem 2.18 (Euler's formula)**.** *Let $P \subset \mathbb{R}^d$ be a d-dimensional convex polytope. Then*

$$\sum_{i=0}^{d-1} (-1)^i f_i(P) = 1 + (-1)^{d-1}.$$

Proof. Lemma 2.9 implies that $I[P] = \sum_F I[\operatorname{relint} F]$, where the summation is taken over all nonempty faces of P, and P itself. Applying the valuation χ to both sides of this equation, the statement follows from Lemma 2.17. □

Exercise 2.5. Consider the five Platonic solids (tetrahedron, cube, octahedron, dodecahedron, icosahedron), and for each of them, let W be the union of its edges. Find the value of $\chi(W)$.

Exercise 2.6. A *Goldberg polyhedron* is a convex 3-dimensional polytope such that each facet is either a pentagon or a hexagon, three facets meet at each vertex, and the polytope has rotational icosahedral symmetry. What is the number of pentagonal facets?

2.3 Face lattice

Recall that a partially ordered set is a set equipped with a binary relation that is reflexive, transitive, and antisymmetric. In the following we consider the partially ordered set of faces of a convex polytope, ordered by inclusion.

Definition 2.19. Let $P \subset \mathbb{R}^d$ be a d-dimensional convex polytope. We define the partially ordered set $\mathcal{F}(P)$ as the set of the faces of P ordered by inclusion.

It turns out that $\mathcal{F}(P)$ has a number of special order-theoretic properties, which we first define in the context of general partially ordered sets.

Definition 2.20. Let $(A, \leq)$ be a partially ordered set.

(i) The *join* (or *supremum*) of a subset $S \subseteq A$ is an element $j \in A$ such that $a \leq j$ for all $a \in S$ and if $j' \in A$ satisfies $a \leq j'$ for all $a \in S$, then also $j \leq j'$.

(ii) The *meet* (or *infimum*) of a subset $S \subseteq A$ is an element $m \in A$ such that $m \leq a$ for all $a \in S$ and if $m' \in A$ satisfies $m' \leq a$ for all $a \in S$, then also $m' \leq m$.

(iii) $(A, \leq)$ is a *lattice* (in the algebraic sense) if every subset of at most two elements has a join and a meet.

(iv) The *opposite* partially ordered set $(A, \leq)^{\mathrm{op}}$ is the same underlying set equipped with the reversed order relation, i.e., $a \leq^{\mathrm{op}} b$ if and only if $b \leq a$.

In a lattice, every finite subset has a join and a meet, and they are unique. The join and the meet operations are dual in the sense that the join of a subset in the opposite partially ordered set is the meet of the same subset in the original partially ordered set. It follows that the opposite of a lattice is a lattice as well.

Definition 2.21. Let $(A, \leq)$ be a partially ordered set.

(i) If $a, b \in A$, then we say that b *covers* a if $a \leq b$, $a \neq b$, and for every $c \in A$ satisfying $a \leq c \leq b$, we have either $a = c$ or $b = c$.

(ii) $\rho : A \to \mathbb{N}$ is a *rank function* if for all $a, b \in A$, $a \leq b$ implies $\rho(a) \leq \rho(b)$, and if b covers a, then $\rho(a) + 1 = \rho(b)$. A partially ordered set equipped with a rank function is called *graded*.

(iii) A *minimum* is an element $0 \in A$ such that for all $a \in A$ the relation $0 \leq a$ holds.

(iv) A *maximum* is an element $1 \in A$ such that for all $a \in A$ the relation $a \leq 1$ holds.

(v) Suppose that a minimum in $(A, \leq)$ exists. An element $a \in A$ is an *atom* if a covers 0.

(vi) Suppose that $(A, \leq)$ is a lattice and has a minimum. $(A, \leq)$ is *atomic* if for every $b \in A \setminus \{0\}$ there exists an atom a such that $a \leq b$.

(vii) Suppose that $(A, \leq)$ is a lattice and has a minimum. $(A, \leq)$ is *atomistic* if every element is the supremum of some atoms in A.

A bounded (one having both a minimum and a maximum) partially ordered set has a rank function if and only if all maximal chains (totally ordered subsets) have the same size. The rank function is unique if we require $\rho(0) = 0$.

Theorem 2.22. *For a convex polytope P, consider the partially ordered set $\mathcal{F}(P)$. Then*

(i) $\mathcal{F}(P)$ is a lattice.

(ii) $\mathcal{F}(P)$ is atomic and atomistic. Its atoms are the vertices of P.

(iii) $\mathcal{F}(P)$ is bounded. The minimum is $\emptyset$ and the maximum is P.

(iv) The function $\rho : \mathcal{F}(P) \to \mathbb{N}$ given by $\rho(F) = \dim(F) + 1$ is a rank function satisfying $\rho(\emptyset) = 0$ (recall the convention $\dim(\emptyset) = -1$).

We call $\mathcal{F}(P)$ the *face lattice* of P.

Example 2.23. *Let V be the set of vertices of a d-dimensional simplex. Then $|V| = d+1$ and for each subset S of V there is a unique face F such that $F \cap V = S$. This provides an order-preserving bijection between the face lattice and the Boolean lattice of subsets of V.*

Definition 2.24. Let $P, Q \subseteq \mathbb{R}^d$ be polytopes. We say that

(i) P and Q are *combinatorially equivalent* if their face lattices are isomorphic,

(ii) P and Q are *dual polytopes* if $\mathcal{F}(P)^{\mathrm{op}}$ is isomorphic to $\mathcal{F}(Q)$.

Thus, the study of face lattices of polytopes is the study of polytopes up to combinatorial equivalence.

Theorem 2.25. *Every polytope has a dual polytope.*

In particular, $\mathcal{F}(P)^{\mathrm{op}}$ is also atomic and atomistic for every polytope P, i.e., every face F is the intersection of the facets that contain F.

Next we study the size of the face lattice. As a graded partially ordered set, we can refine this to the sizes of *rank levels*. Recall that $f_j(P)$ denotes the number of j-faces of the convex polytope P.

Definition 2.26. The *f-vector* of the convex polytope P is the vector

$$f(P) = (f_0(P), f_1(P), \dots, f_d(P)).$$

We will be interested in constraints on $f(P)$ depending on $f_0(P) =: n$, the number of vertices, in particular the maximum of $f_j(P)$ over all polytopes with n vertices. For $d = 2$, P is a convex polygon, therefore $f(P) = (n, n, 1)$ for some $n \geq 3$. For $d = 3$, it is known that $f_1(P) \leq 3n - 6$ and $f_2(P) \leq 2n - 4$, and equality holds if every face is a triangle (use $3f_2(P) \leq 2f_1(P)$ and Euler's formula $2 = f_0(P) - f_1(P) + f_2(P)$). However, when $d \geq 4$, the number of faces can be superlinear in n.

Exercise 2.7. Describe the face lattice of the d-dimensional simplex.

Exercise 2.8. Describe the face lattice of the d-dimensional cube.

2.4 Simple and simplicial polytopes

We introduce a special class of polytopes.

Definition 2.27. A d-dimensional polytope is *simple* if every vertex is incident with exactly d facets. A polytope is *simplicial* if every facet is a simplex.

Example 2.28.

(i) Every convex polygon is both simple and simplicial.

(ii) *The (3-dimensional) cube and the dodecahedron are simple polytopes, whereas the octahedron and the icosahedron are simplicial. A tetrahedron is both simple and simplicial.*

(iii) *The regular* 120*-cell is a simple* 4*-dimensional polytope and the regular* 600*-cell is a simplicial* 4*-dimensional polytope. The regular* 24*-cell is neither simple nor simplicial.*

(iv) *For every d, the d-dimensional cube is simple, the cross polytope is simplicial, and the d-simplex is both simple and simplicial.*

A polytope is simple if and only if its dual is simplicial. If the vertices of P are in general position, then P is simplicial.

The following lemma reduces the study of upper bounds on f-vectors to simplicial or simple polytopes.

Lemma 2.29. *Let P be a d-dimensional convex polytope.*

(i) *There exists a d-dimensional simplicial polytope Q such that $f_0(Q) = f_0(P)$ and $f(Q) \geq f(P)$ (componentwise).*

(ii) *There exists a d-dimensional simple polytope Q such that $f_{d-1}(Q) = f_{d-1}(P)$ and $f(Q) \geq f(P)$.*

Theorem 2.30. *Let P be a simple d-dimensional polytope. Then the followings hold.*

(i) *Every vertex of P has exactly d neighbours.*

(ii) *If v is a vertex of P and $\{u_1, u_2, \dots, u_k\}$ is a k-element subset of its neighbours for some $k \leq d$, there is a unique k-face of P that contains $v, u_1, \dots, u_k$ but no other neighbour of v.*

(iii) *The intersection of any $k \leq d$ facets containing v is a $(d-k)$-face of P.*

(iv) *Let $l : \mathbb{R}^d \to \mathbb{R}$ be linear such that $l(u_i) \leq l(v)$ holds for all neighbours $u_1, \dots, u_d$ of v. Then $l|_P$ is maximal at v.*

(v) *Every face of P is a simple polytope.*

Proof. Applying a suitable affine bijection, we may assume that v is the origin, the d facets incident with it are contained in the coordinate hyperplanes, and P is contained in the nonnegative orthant.

(i): Regarding P as an H-polytope, in addition to the half spaces $H_i = \{x \in \mathbb{R}^d | x_i \geq 0\}$, the remaining defining half spaces contain the origin in their interiors. This implies that the edges that are incident with v are subsets of the positive halves of the coordinate axes.

(ii): Let $i_1, i_2, \dots, i_{d-k}$ be the indices of the coordinates that are zero at $u_1, \dots, u_k$. Then the hyperplane $x_{i_1} + x_{i_2} + \dots + x_{i_{d-k}} = 0$ intersects P in a k-face that contains $v, u_1, \dots, u_k$ and no other neighbour of v.

(iii): The intersection of k coordinate hyperplanes is an affine subspace of dimension $d-k$, and it intersects P in a $d-k$-dimensional face.

(iv): Since P is a subset of the nonnegative orthant, which is the cone generated by the edge vectors starting at v, every point is a conic combination of the neighbours $u_1, \dots, u_d$. By assumption, $l(u_i) \geq 0$ for all $i = 1, \dots, d$, therefore for any point x with nonnegative coordinates the inequality $l(x) \leq l(v)$ holds.

(v): If F is a facet that contains v, then it is the intersection of a coordinate hyperplane with P, and the intersections of F with the remaining $d-1$ hyperplanes are the facets of F that contain v.

□

2.5 h-vector

For the study of f-vectors, it is useful to introduce another, related vector.

Definition 2.31. Let $P \subseteq \mathbb{R}^d$ be a convex polytope and $l : \mathbb{R}^d \to \mathbb{R}$ a linear functional such that $l(u) \neq l(v)$ whenever $[u, v]$ is an edge of P (i.e., when u and v are *neighbours*). Let us say that u is *below* v (or v is *above* u) if $l(u) \leq l(v)$. u is an *upper neighbour* of v if it is a neighbour and it is above v. Likewise, a u is a *lower neighbour* of v if it is a neighbour and it is below v.

For a vertex v of P, the *index* with respect to l is the number of lower neighbours of v. We denote by $h_k(P, l)$ denote the number of vertices of P having index k with respect to l.

If P is a simple d-dimensional polytope and l is as before, then we introduce the vector $h(P, l) = (h_0(P, l), h_1(P, l), \dots, h_d(P, l))$.

If P is a simple d-dimensional polytope, then

$$\begin{aligned} h(P, -l) &= (h_0(P, -l), h_1(P, -l), \dots, h_d(P, -l)) \\ &= (h_d(P, l), h_{d-1}(P, l), \dots, h_0(P, l)), \end{aligned}$$

since each of the d neighbours of a vertex v is either above or below v.

Theorem 2.32. *Let P be a simple d-dimensional polytope and $l : \mathbb{R}^d \to \mathbb{R}$ a linear functional such that $l(u) \neq l(v)$ whenever $[u, v]$ is an edge of P. Then the equality*

$$f_i(P) = \sum_{k=i}^{d} \binom{k}{i} h_k(P, l)$$

holds for all $i = 0, \dots, d$.

Proof. Each face F has a unique highest vertex, since the subset of F where l is maximal is a face that does not contain an edge. Therefore we can count the number of i-faces by adding, for each vertex v, the number of i-faces whose highest vertex is v. Such a face can only contain neighbours of v that are below v.

Let v have index k. For each i-element subset Y of the set of lower neighbours of v, there is a unique i-face F that contains v and Y and does not contain neighbours that are not in Y. v is the highest vertex of this F. In total, there are $\binom{k}{i}$ faces of dimension i whose highest vertex is v. The claim follows by summing over all vertices. □

Theorem 2.33. *Let P be a simple d-dimensional polytope and $l : \mathbb{R}^d \to \mathbb{R}$ a linear functional such that $l(u) \neq l(v)$ whenever $[u, v]$ is an edge of P. Then we have*

$$h_k(P, l) = \sum_{i=k}^{d} (-1)^{i-k} \binom{i}{k} f_i(P).$$

In particular, $h(P, l)$ is independent of l.

Proof. Let us encode both vectors in the generating functions $F_P(t) = \sum_{i=0}^{d} f_i(P)t^i$ and $H_{P,l}(t) = \sum_{k=0}^{d} h_k(P, l)t^k$. By Theorem 2.32, we have

$$\begin{aligned} F_P(t) &= \sum_{i=0}^{d} f_i(P)t^i \\ &= \sum_{i=0}^{d} \sum_{k=i}^{d} \binom{k}{i} h_k(P, l)t^i \\ &= \sum_{k=0}^{d} \sum_{i=0}^{k} \binom{k}{i} h_k(P, l)t^i \\ &= \sum_{k=0}^{d} h_k(P, l) \sum_{i=0}^{k} \binom{k}{i} t^i \\ &= \sum_{k=0}^{d} h_k(P, l)(t + 1)^k \\ &= H_{P,l}(t + 1), \end{aligned}$$

therefore $H_{P,l}(t) = F_P(t - 1)$, which does not depend on l. Since

$$\begin{aligned} \sum_{k=0}^{d} h_k(P, l)t^k &= H_{P,l}(t) \\ &= F_P(t - 1) \\ &= \sum_{i=0}^{d} f_i(P)(t - 1)^i \\ &= \sum_{i=0}^{d} \sum_{k=0}^{i} (-1)^{i-k} \binom{i}{k} f_i(P)t^k \\ &= \sum_{k=0}^{d} \left[\sum_{i=k}^{d} (-1)^{i-k} \binom{i}{k} f_i(P) \right] t^k, \end{aligned}$$

the claim follows by comparing the coefficients. □

From now on we will drop l from the notation and write $h(P) = (h_0(P), h_1(P), \dots, h_d(P))$, and call it the *h-vector* of the simple polytope P.

Corollary 2.34 (Dehn–Sommerville equations)**.** *If P is a d-dimensional simple polytope, then the h-vector satisfies $h_k(P) = h_{d-k}(P)$ for every $k = 0, 1, \dots, d$.*

Remark 2.35. *In terms of the f-vector, the equality $h_0(P) = h_d(P)$ expands to*

$$\sum_{i=0}^{d} (-1)^i \binom{i}{0} f_i(P) = (-1)^{d-d} \binom{d}{d} f_d(P),$$

i.e.,

$$\sum_{i=0}^{d-1} (-1)^i f_i(P) = 1 + (-1)^{d-1},$$

which is the Euler formula (specialized to simple polytopes).

The equality $h_1(P) = h_{d-1}(P)$ says

$$\sum_{i=1}^{d} (-1)^{i-1} i f_i(P) = f_{d-1}(P) - d f_d(P).$$

Exercise 2.9. Find the h-vector of the d-dimensional simplex.

Exercise 2.10. Find the h-vector of the d-dimensional cube.

Exercise 2.11. Let P be a 3-dimensional simple polytope with n vertices. Using the Dehn–Sommerville equations, prove that $f_1(P) = \frac{3}{2}n$ and $f_2(P) = 2 + \frac{n}{2}$.

Exercise 2.12. Let P be a 4-dimensional simple polytope. Prove that $f_1(P) = 2f_0(P)$.

2.6 Upper bound theorem

A facet of a simple polytope is simple by part Part (v) of Theorem 2.30, therefore it also has an h-vector.

Lemma 2.36. *Let $P \subseteq \mathbb{R}^d$ be a d-dimensional simple polytope and F a facet. Then $h(F) \leq h(P)$ (componentwise), and if every set of $k+1$ facets of P has nonempty intersection, then $h_k(F) = h_k(P)$.*

Proof. Let $l_0 : \mathbb{R}^d \to \mathbb{R}$ be a linear functional such that, with $m := \min_{x \in P} l_0(x)$, $x \in P$ and $l_0(x) = m$ implies $x \in F$ (if we picture l_0 as the vertical coordinate, we may imagine that P rests on F). Let $l : \mathbb{R}^d \to \mathbb{R}$ be a small perturbation of l_0 such that l takes different values on any pair of neighbouring vertices of P, and $l(x) < l(y)$ whenever x and y are vertices of P such that $x \in F$ and $y \notin F$.

Let v be a vertex in F, and let its index (as a vertex of F) with respect to l be k. This means that there are k neighbours of v in F below v and $d-1-k$ neighbours above v. In P, v has an additional neighbour, which is above v by the choice of l. Therefore the index of v as a vertex of P is also k. Since h_k counts the number of index-k vertices, we have $h_k(F) \leq h_k(P)$.

Suppose that every set of $k+1$ facets of P have nonempty intersection. Let v be a vertex of P of index k. Let $u_1, \ldots, u_k$ be its lower neighbours and $u_{k+1}, u_{k+2}, \ldots, u_d$ its upper neighbours. There is a $d-k$-face G that contains $v, u_{k+1}, u_{k+2}, \ldots, u_d$ by Part (ii) of Theorem 2.30, and by Part (iv), l is minimal on G at v. By Part (iii) of Theorem 2.30, G is the intersection of k facets $F_1, \ldots, F_k$ where $u_i \notin F_i$ but v and the other lower neighbours are contained in F_i. Then $F \cap G = F \cap F_1 \cap \cdots \cap F_k$ is the intersection of $k+1$ facets, therefore it is nonempty by assumption. Let one of the vertices of $F \cap G$ be w. Since l is minimal on G at v, we have $l(v) \leq l(w)$. By the choice of l, this implies $v \in F$. We conclude that every vertex of P of index k is in F, therefore $h_k(F) \geq h_k(P)$. □

Proposition 2.37. *Let P be a simple d-dimensional polytope. Then for $k = 0, 1, \ldots, d-1$ we have*

$$\sum_{\substack{F \subseteq P \\ \text{is a facet}}} h_k(F) = (d-k)h_k(P) + (k+1)h_{k+1}(P).$$

Proof. Let $l : \mathbb{R}^d \to \mathbb{R}$ be a linear functional that takes different values on neighbouring vertices of P (and hence of every facet).

Let v be a vertex. There are exactly d faces that contain v, intersecting the set of the d neighbours of v in all possible subsets of size $d-1$. If the omitted neighbour is above v, then the index of v in as a vertex of P is the same as the index as a vertex of the facet. Otherwise, if the omitted neighbour is below v, then the index as a vertex of P is one plus the index as a vertex of the facet.

Thus, a vertex v can contribute to the left hand side in two ways: either it has index k as a vertex of P, and the facet is obtained by omitting an upper neighbour ($d-k$ possibilities), or it has index $k+1$ and the facet is obtained by omitting a lower neighbour ($k+1$ possibilities). □

Theorem 2.38. *Let P be a simple d-dimensional polytope and let $n = f_{d-1}(P)$. Then for $k = 0, 1, \ldots, d$ the inequality*

$$h_k(P) \leq \min\left\{ \binom{n-d+k-1}{k}, \binom{n-k-1}{d-k} \right\}$$

holds.

If the intersection of any set of k facets of P is nonempty, then equality holds.

Proof. By Lemma 2.36 and Proposition 2.37, we have

$$(d-i)h_i(P) + (i+1)h_{i+1}(P) = \sum_{\substack{F \subseteq P \\ \text{is a facet}}} h_i(F) \leq n h_i(P)$$

for all i, which can be rearranged as

$$h_{i+1}(P) \leq \frac{n-d+i}{i+1} h_i(P),$$

with equality if the intersection of any set of i facets of P is nonempty. Since $h_0(P) = 1$, this implies

$$h_k(P) \leq \prod_{i=0}^{k-1} \frac{n-d+i}{i+1} = \binom{n-d+k-1}{k},$$

with equality if the intersection of any set of k facets of P is nonempty.

By the Dehn–Sommerville equations, we also have

$$h_k(P) = h_{d-k}(P) \leq \binom{n-d+(d-k)-1}{d-k} = \binom{n-k-1}{d-k},$$

which concludes the proof. □

Theorem 2.39 (Upper Bound Theorem)**.** *Let P be a d-dimensional polytope with n facets. If $i \leq \lfloor \frac{d}{2} \rfloor$, then*

$$f_i(P) \leq \sum_{k=i}^{\lfloor d/2 \rfloor} \binom{k}{i}\binom{n-d+k-1}{k} + \sum_{k=\lfloor d/2 \rfloor+1}^{d} \binom{k}{i}\binom{n-k-1}{d-k},$$

while if $i > \lfloor \frac{d}{2} \rfloor$, then

$$f_i(P) \leq \binom{n}{d-i}.$$

Proof. We prove the bound for $i > \lfloor \frac{d}{2} \rfloor$ first. Since every i-face is the intersection of $d-i$ facets, $f_i(P)$ is at most the number of $d-i$-element subsets of the set of facets.

Next we prove the $i \leq \lfloor \frac{d}{2} \rfloor$ case. By Lemma 2.29, there is a simple d-dimensional polytope Q with n facets and $f(P) \geq f(Q)$. We use Theorem 2.32 and Theorem 2.38 to estimate the number of faces of Q:

$$\begin{aligned} f_i(P) &\leq f_i(Q) \\ &= \sum_{k=i}^{d} \binom{k}{i} h_k(Q) \\ &\leq \sum_{k=i}^{d} \binom{k}{i} \min\left\{\binom{n-d+k-1}{k}, \binom{n-k-1}{d-k}\right\} \\ &\leq \sum_{k=i}^{\lfloor d/2 \rfloor} \binom{k}{i}\binom{n-d+k-1}{k} + \sum_{k=\lfloor d/2 \rfloor+1}^{d} \binom{k}{i}\binom{n-k-1}{d-k}. \end{aligned}$$

□

In the following we show that the upper bounds in Theorem 2.39 cannot be improved. In fact, we will construct polytopes depending on d and n that saturate every inequality.

Definition 2.40. The *moment curve* in $\mathbb{R}^d$ is the curve $\gamma : \mathbb{R} \to \mathbb{R}^d$, $\gamma(t) = (t, t^2, \dots, t^d)$.
The convex hull of a finite set of points on γ is a *cyclic polytope*.

Note that any $d+1$ points on the moment curve are affinely independent, therefore a hyperplane intersects it in at most d points. It follows that cyclic polytopes are simplicial.

Proposition 2.41. *Let $C(d,n) = \operatorname{conv}(\{v_1, \dots, v_n\})$ be a cyclic polytope with $v_i = \gamma(t_i)$, where $t_1, \dots, t_n$ are distinct real numbers. Let $I \subseteq [n]$ be a set such that $|I| \le \frac{d}{2}$. Then $F = \operatorname{conv}(\{v_i\}_{i \in I})$ is a face of $C(d,n)$.*

Proof. Let $I \subseteq [n]$, $|I| \le \frac{d}{2}$. Consider the nonnegative polynomial

$$p(\tau) = \prod_{i \in I} (\tau - t_i)^2 = \sum_{m=0}^{d} a_m \tau^m,$$

and the linear functional $l(x_1, \dots, x_d) = \sum_{m=1}^{d} a_m x_m$. For the points of the moment curve, we have $l(t, t^2, \dots, t^d) = p(t) - a_0 \ge -a_0$. If $i \in I$, then $p(t_i) = 0$, therefore equality holds. This implies that the hyperplane $H = \left\{ x \in \mathbb{R}^d \middle| l(x) = -a_0 \right\}$ intersects $C(d,n)$ in F, and $C(d,n)$ is contained in one of the closed half spaces whose boundary is H. □

Proposition 2.42. *The dual of $C(d,n)$ saturates the inequalities in Theorems 2.38 and 2.39.*

Proof. Let P be a dual of $C(d,n)$. By Theorem 2.32 and the Dehn–Sommerville equations, for every i the number $f_i(P)$ is a conic combination of $h_0(P), \dots, h_{\lfloor d/2 \lfloor}(P)$. Therefore if the upper bounds in Theorem 2.38 hold with equalities, then also the bound in Theorem 2.39 is saturated.

For $k \le \frac{d}{2}$, the intersection of any k facets of P is a $(d-k)$-face of P. This implies that the inequalities in Theorem 2.38 hold with equalities. □

To summarize, the dual of $C(d,n)$ maximizes the number of i-faces among d-dimensional convex polytopes with n facets, for each $0 \le i \le d \le n-1$. By taking duals, $C(d,n)$ maximizes the number of i-faces among d-dimensional convex polytopes with n vertices.

2.7 Graphs of polytopes

Following the usual notation and terminology, we write a simple and undirected graph as $G = (V, E)$ where V is a set and $E \subseteq \binom{V}{2}$, and the elements of V and E are called vertices and edges, respectively. For a subset $S \subseteq V$ we denote by $G - S$ the graph with vertex set $V \setminus S$ and edge set $E \cap \binom{V \setminus S}{2}$ (i.e., the vertices in S along with any edge incident with them removed).

Definition 2.43. A *walk* in a graph G is a sequence $v_1, v_2, \ldots, v_n$ of vertices such that $\{v_i, v_{i+1}\}$ is an edge for all $i = 1, 2, \ldots, n-1$. The graph $G = (V, E)$ is *connected* if for all $u, v \in V$ there is a walk starting at u and ending at v.

A graph G is *k-connected* (or *k-vertex-connected*) if it has at least $k+1$ vertices and for all subsets S of $k-1$ vertices, $G - S$ is connected.

Definition 2.44. Let P be a polytope. The *graph of P* is the graph $G(P)$ whose vertex set is the set of vertices of P, and $\{u, v\}$ is an edge of $G(P)$ if $[u, v]$ is an edge of P.

Theorem 2.45 (Balinski)**.** *Let P be a d-dimensional polytope. Then $G(P)$ is d-connected.*

Proof. Since P is d-dimensional, $G(P)$ it has at least $d+1$ vertices. Let $S \subseteq V$, $|S| = d-1$. We need to show that $G(P) - S$ is connected. Let $s = \sum_{v \in S} \frac{1}{d-1} v$, and let F be the unique face that contains s in its relative interior.

If $F \neq P$, then it is the intersection of P with a supporting hyperplane, therefore $S \subseteq F$ as well. Let $l : \mathbb{R}^d \to \mathbb{R}$ be a linear functional such that for $\max_{x \in P} l(x) = M$ the equality $F = l^{-1}(M) \cap P$ holds. Let $m = \min_{x \in P} l(x)$ and $F_0 = l^{-1}(m) \cap P$. Every vertex v not contained in F_0 has at least one neighbour u such that $l(u) < l(v)$. This implies that for any $v \in V \setminus S$ there exists a walk $v = v_1, v_2, \ldots, v_n \in F_0$. By induction on the dimension of the polytope, $G(F_0)$ is a connected subgraph of $G(P) - S$, therefore $G(P) - S$ is also connected.

Suppose now that $F = P$. Let l be a nonzero linear functional that takes a constant value c on S and at least one vertex v_0 in $V \setminus S$. Let $F_{\min}$ and $F_{\max}$ be the faces of P where l is minimal and maximal. If a vertex v satisfies $l(v) \leq c$, then there is a path in $G(P) - S$ starting at v and ending in $F_{\min}$ (as in the previous case, by appending a neighbour where l strictly decreases). Similarly, if $l(v) \geq c$, then there is a path in $G(P) - S$ starting at v and ending in $F_{\max}$. But v_0 is connected to both $F_{\min}$ and $F_{\max}$, therefore $G(P) - S$ is connected. □

Our next goal is to prove Steinitz' theorem: simple, 3-connected, planar graphs are graphs of 3-polytopes. The main idea in Steinitz' proof is to show that such a graph can be built from K_4 with a sequence of transformations (called ΔY operations) that preserve realizability as the graph of a 3-polytope. The graph-theoretic part was strengthened by Epifanov, who showed that by allowing series and parallel reductions, planar graphs can be reduced to a single edge between a specified pair of vertices. We follow a simplified proof of this result due to Truemper, with the following steps:

- every simple, 3-connected, planar graph is a minor of a suitably large grid graph
- if a 3-connected graph can be transformed by ΔY reductions to K_4, then so is every 3-connected minor
- grid graphs can be transformed into K_4 by ΔY reductions
- if a graph G can be transformed by a ΔY reduction into another graph G' that is the graph of a polytope, then G is also the graph of a polytope.

In the following we also need to allow graphs to be non-simple. We continue to use the notation $G = (V, E)$, but the edge set E can no longer be identified with a subset of the set of unordered pairs of vertices. Instead, an relation called *incidence* is assumed to be given between V and E such that every $e \in E$ is incident with exactly 1 or 2 vertices. An edge that is incident with one vertex is a *loop*. Edges that are incident with the same set of vertices are called *parallel*.

Definition 2.46. A *drawing* of a graph $G = (V, E)$ is a map $d : V \to \mathbb{R}^2$ and for each $e \in E$ a contiuous injective map $\gamma_e : [0, 1] \to \mathbb{R}^2$ such that

(i) d is injective,

(ii) for each edge e, if $u, v \in V$ are incident with e (where $u = v$ is allowed), then $\gamma_e(\{0, 1\}) = \{d(u), d(v)\}$,

(iii) for each edge e, $\gamma_e((0, 1)) \cap d(V) = \emptyset$,

A *crossing* of a drawing is a point $x \in \mathbb{R}^2$ such that there exist distinct edges e, f such that $x \in \gamma_e((0, 1)) \cap \gamma_f((0, 1))$.

A graph G is *planar* if it has a drawing without crossing.

We consider the following local transformations of graphs.

Definition 2.47. The *deletion* of an edge removes it from E without changing V. The *contraction* of an edge e that is not a loop removes it from E and replaces the vertices u, v incident with e with a new vertex that is incident with precisely those edges that were previously incident with either u of v. Graphs that can be obtained from G by a sequence of deletions and contractions is a *minor* of G. The reverse operation of a contraction is called *splitting*.

An *SP reduction* is a sequence of deletions of edges that are parallel with others and contractions of edges that are incident with at least one degree-2 vertex (*series edges*).

A *Δ-to-Y operation* replaces a nonseparating triangle with a 3-star that connects the same vertices. The reverse is called a *Y-to-Δ operation.* Δ-to-Y operations and Y-to-Δ operations are collectively referred to as *ΔY operations.* A *simple ΔY reduction* is any ΔY operation followed by all possible SP reductions.

In the following, we use duality of polytopes (that can be realized by the polar construction) to reduce the number of cases that we need to consider. The following notions correspond to each other under duality:

$$
\begin{array}{ccc}
\text{planar graph } G & \longleftrightarrow & \text{dual graph } G^*, \\
\text{contracting series edges} & \longleftrightarrow & \text{deleting parallel edges}, \\
k\text{-connected} & \longleftrightarrow & k\text{-connected}, \\
\text{nonseparating triangle} & \longleftrightarrow & \text{3-star}, \\
\Delta\text{-to-}Y \text{ transformation} & \longleftrightarrow & Y\text{-to-}\Delta \text{ transformation.}
\end{array}
$$

Lemma 2.48. *Let G be a graph and v a degree-3 vertex with three different neighbours, and let G' be the graph obtained by applying a Y-to-Δ operation at v.*

(i) If G is a 2-connected graph, then G' is also 2-connected.

(ii) If G is 3-connected, $G \neq K_4$, then, after deleting any parallel edges of G', the resulting graph is also 3-connected.

Proof. Any set of one or two separating vertices in G' is also a separating set in G. This implies that if we delete parallel edges, then 3-connectedness is preserved. □

Lemma 2.49. *Let G be a 3-connected planar graph and G' the result of a simple ΔY reduction. If G' is the graph of a 3-dimensional polytope then G is also the graph of a 3-dimensional polytope.*

Proof. By duality, it is sufficient to consider a Δ-to-Y transformation (followed by all possible SP reductions). Given a polytope P' with graph G', we can obtain a polytope P with graph G by cutting P' near the degree-3 vertex corresponding to the center of the "Y". We can distinguish four cases depending on the number of degree-3 vertices of the triangle, which correspond to different ways of cutting (see Figure 2). □

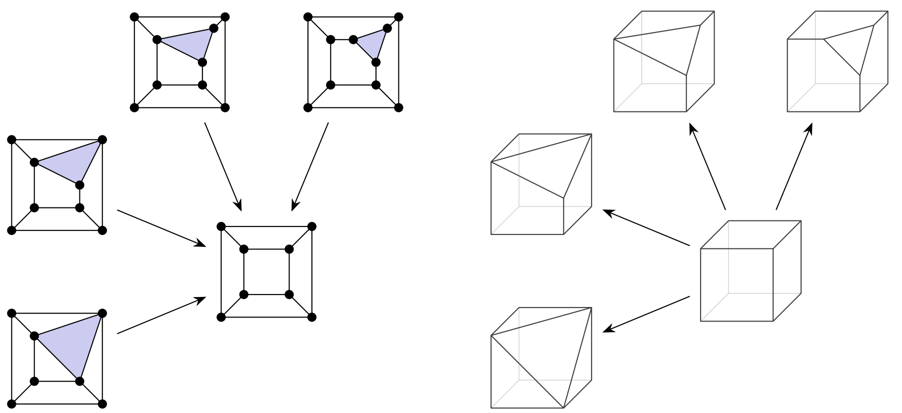

Figure 2: The four kinds of cuts of a polytope P' near a degree-3 correspond to simple ΔY reductions from different graphs to $G(P')$, distinguished by the number of degree-3 vertices of the triangle.

Let us call a 2-connected graph ΔY*-reducible* if it can be transformed into C_2 (two vertices with two parallel edges) by a sequence of ΔY transformations and SP reductions.

Lemma 2.50. *Let G be a planar graph that is ΔY-reducible, and let H be a 2-connected minor. Then H is ΔY-reducible.*

Proof. If H has series or parallel edges, then we can apply SP reductions to remove them both in G and in H. We may therefore assume that H has no series or parallel edges.

We use induction on the number of reduction steps for G. If the reduction for G starts with a series or parallel reduction, then H is a minor of the reduced graph as well, since it does not contain both edges.

If the reduction of G starts with a ΔY operation, then we may assume that this is a Δ-to-Y operation by considering duals if necessary, resulting in a graph G'.

If all three edges in the triangle are present in H, then the same Δ-to-Y operation can be applied to H as well, resulting in a minor H' of G'. By induction, H' is ΔY-reducible, and therefore H is ΔY-reducible as well.

Otherwise some of them are deleted or all of them are contracted when H was formed. In the latter case, we may equivalently delete first one edge and then contract the remaining two. Thus we may assume that the first step of obtaining H from G is the deletion of an edge. But then we get the same graph by contracting the opposite edge in G', therefore H is a minor of G' as well. By induction, H is ΔY-reducible. □

For $m, n \in \mathbb{N}$, the *grid graph* $G(m, n)$ is the graph with vertex set $([1, m] \times [1, n]) \cap \mathbb{Z}^2$, where (i, j) and (i', j') are neighbours if and only if $|i - i'| + |j - j'| = 1$.

Lemma 2.51. *Let G be a planar graph. Then G is a minor of a grid graph.*

Proof. By repeatedly splitting the vertices of G we can obtain a graph G' such that every vertex has degree at most 3. Then G' is still planar and G is a minor of G'. We then take any drawing of G', replace the edges by simple polygonal chains consisting of axis-parallel line segments, then modify them in such a way that every vertex has rational coordinates. After clearing denominators and subdividing at each integer point, we obtain a subgraph of a suitably large grid. □

Lemma 2.52. *Let $m, n \geq 3$. Then the grid graph $G(m, n)$ is ΔY-reducible to K_4.*

Proof. If there is an edge e between two neighbours of a degree-3 vertex v, then the deletion of e is a simple ΔY reduction: we apply a Δ-to-Y operation to the triangle formed by v and the endpoints of e, then contract v with the new vertex.

If there is an edge e between two neighbours of a degree-4 vertex v, then we remove the edge and add another between the other two neighbours by a Δ-to-Y operation at this triangle and a Y-to-Δ operation at v.

Let $G'(m, n)$ be the graph obtained from $G(m, n)$ by a series reduction at the $(1, 1)$ vertex (so $G(m, n)$ is ΔY-reducible to $G'(m, n)$). We apply series reduction at the opposite corner (m, n), and move the new edge by a sequence of the degree-4 moves until it hits the left column or bottom row. At this point the edge can be deleted by a degree-3 move or it is a parallel edge that can be removed.

We continue with the $(m - 1, n)$ vertex in a similar way, applying a series reduction, a sequence of degree-4 moves and deletion of the edge when it arrives in the left column or bottom row. After repeating this process with the vertices $(m - 2, n), \ldots, (3, n)$, The last square is removed by two series and one parallal reduction.

In this way we see that $G'(m, n)$ is ΔY-reducible to $G'(m, n - 1)$ if $n \geq 4$. Since $G'(m, n)$ is isomorphic to $G'(n, m)$, we obtain that $G'(m, n)$ is ΔY-reducible to $G'(3, 3)$ as well.

From $G'(3, 3)$ we proceed by series reducing the remaining 3 corners, applying a Δ-to-Y operation and two series reductions. □

Corollary 2.53. *Let G be a 3-connected planar graph. Then G can be reduced to K_4 by a sequence of simple ΔY-reductions.*

Proof. We use induction on the number of edges. By Lemmas 2.50 to 2.52, G is ΔY-reducible. We apply the corresponding moves until a parallel or series edge is created. After an SP reduction, we obtain a 3-connected planar graph (Lemma 2.48) with fewer edges. □

Theorem 2.54 (Steinitz). *A finite, simple graph G is the graph of a 3-dimensional polytope if and only if it is 3-connected and planar.*

Proof. Let P be a 3-dimensional polytope. By Balinski's theorem, the graph $G = G(P)$ is 3-connected. If $0 \in \operatorname{int}(P)$ and the z-axis contains no vertex, then the map

$$(x, y, z) \mapsto \left(\frac{x}{\sqrt{x^2 + y^2 + z^2} - z}, \frac{y}{\sqrt{x^2 + y^2 + z^2} - z} \right)$$

determines a drawing (by applying to the vertices and composing with the parameterized segments corresponding to the edges). Therefore $G(P)$ is planar.

Now assume that G is a 3-connected, simple, planar graph. If $G = K_4$, then $G = G(P)$ where P is a tetrahedron. Otherwise, G can be reduced to K_4 by a sequence of ΔY reductions by Corollary 2.53. Using Lemma 2.49 at each step from the end of the reduction, we can see that G is also the graph of a 3-polytope. □

Remark 2.55. *A 3-connected, simple, planar graph has a drawing in which every edge is drawn as a straight line segment and every face (including the outer face) is a convex polygon. This follows from the Steinitz theorem by noting that such a graph is the graph of a convex 3-dimensional polytope P, and a drawing can be obtained by projecting the vertices and edges of P from a point x close to a facet F to the plane containing F.*

Exercise 2.13. A convex polytope P is *dimensionally ambiguous* if there is a convex polytope Q such that $\dim P \neq \dim Q$ and $G(P) \simeq G(Q)$. Show that a d-simplex is dimensionally ambiguous if and only if $d \geq 5$.

Exercise 2.14. Let P be a convex 3-dimensional polytope, and let v_3 denote the number of degree-3 vertices and f_3 the number of triangular faces. Show that $v_3 + f_3 \geq 8$.

Exercise 2.15. Find a ΔY reduction of the graph of the cube.

2.8 Convex hull computation in the plane

In this secion we briefly discuss the computation of the convex hull. The general task is, given a finite set $V \subseteq \mathbb{R}^d$ of points as input, to output a description of the convex hull. There are various ways kinds of descriptions one might aim for: give a list of half spaces whose intersection is $\operatorname{conv}(V)$, find all vertices and faces (given by the sets of their vertices), the face lattice, or the incidence relation between vertices and facets. In a different direction, one might wish to test membership in a convex hull or draw random samples from the uniform distribution on $\operatorname{conv}(V)$.

We consider only the planar case, where the task can be formulated as listing the vertices of the convex hull of V in (say) counterclockwise order. We assume that the points are given by pairs of coordinates with respect to a given Cartesian coordinate system. The steps of the *monotone chain algorithm* (also known as *Andrew's algorithm*) are the following:

1. Sort the points lexicographically by their x and y coordinates to get a list P.
2. Let U and L be empty lists.
3. For $i = 1, 2, \dots, n$:
   - While L contains at least two points and the last two points and $P[i]$ do not form a counterclockwise turn: remove the last point in L.
   - Append $P[i]$ to L.
4. For $i = n, n-1, \dots, 1$:
   - While U contains at least two points and the last two points and $P[i]$ do not form a counterclockwise turn: remove the last point in U.
   - Append $P[i]$ to U.
5. Remove the last point of L and U.
6. Return the concatenation of L and U.

Remark 2.56. *The edges determined by L and U before removing their last points are the* lower *and* upper hulls *of V.*

Theorem 2.57. *Given a list of n points in $\mathbb{R}^2$, the monotone chain algorithm returns the vertices of their convex hull in $O(n \log n)$ time.*

Proof. Let V be the set of input points and let H be the polygon formed by the vertices in the output (the edges are formed by adjacent elements in the list). Since every vertex of P is in V, we have $H \subseteq \mathrm{conv}(V)$.

Consider the sets

$$U^{\downarrow} = \left\{(x, y) \in \mathbb{R}^2 \middle| \exists y' \in \mathbb{R} : y \leq y' \text{ and } (x, y') \in \mathrm{conv}(U)\right\}$$

and

$$L^{\uparrow} = \left\{(x, y) \in \mathbb{R}^2 \middle| \exists y' \in \mathbb{R} : y' \leq y \text{ and } (x, y') \in \mathrm{conv}(L)\right\}.$$

By construction, both $U^{\downarrow}$ and $L^{\uparrow}$ are convex and can only get larger in every step, and after step i they contain $P[i]$. Therefore, when the algorithm terminates, $\mathrm{conv}(V) \subseteq U^{\downarrow} \cap L^{\uparrow} = H$.

The statement on the runtime follows from the fact that sorting is possible in $O(n \log n)$ time, and the subsequent loops require $O(n)$ steps (each point is added exactly once and removed at most once). □

Remark 2.58. *The upper bound on the runtime as a function on the input size is optimal: it is known that sorting a list of numbers requires $\Omega(n \log n)$ comparisons (worst case) and sorting $x_1, x_2, \ldots, x_n$ is equivalent to finding the convex hull of the points $(x_1, x_1^2), (x_2, x_2^2), \ldots, (x_n, x_n^2)$.*

3 Incidence problems

In the following we will use standard notation in connection with the asymptotic behaviour of positive sequences, which we first recall (computer science definitions).

Definition 3.1. Let $f, g : \mathbb{N} \to \mathbb{R}_{\geq 0}$, $g(n) > 0$ for all sufficienty large n. We write

(i) $f = O(g)$ if $\limsup_{n\to\infty} \dfrac{f(n)}{g(n)} < \infty$,

(ii) $f = o(g)$ if $\limsup_{n\to\infty} \dfrac{f(n)}{g(n)} = 0$,

(iii) $f = \Omega(g)$ if $\liminf_{n\to\infty} \dfrac{f(n)}{g(n)} > 0$,

(iv) $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$.

We note that some authors prefer notations like $f \in O(g)$, etc., which may be considered more appropriate as it emphasizes that the relation is not symmetric. Nevertheless, the symbol $=$ in this sense is in widespread use.

3.1 Point–line incidences

We study the maximum possible number of incidences between points and lines.

Definition 3.2. Let P be a set of points and L a set of lines in the plane. The number of *incidences* is $I(P, L) = |\{(p, l) \in P \times L | p \in l\}|$.

For $m, n \in \mathbb{N}$, we denote by $I(m, n)$ the maximum of $I(P, L)$ over all sets of points and lines satisfying $|P| = m$, $|L| = n$.

Since $I(P, L)$ is the cardinality of a subset of $P \times L$, we have $I(P, L) \leq |P| \cdot |L|$ and therefore $I(m, n) \leq mn$.

We will not determine exact values but aim for asymptotically optimal lower and upper bounds.

Proposition 3.3. $I(n, n) \in \Omega(n^{4/3})$.

Proof. Let $n = 4k^3$ where $k \in \mathbb{N}$. Let $l_{a,b} = \{(x, y) \in \mathbb{R}^2 | y = ax + b\}$, and consider the sets

$$P = \{(i, j) \in \mathbb{Z}^2 | 0 \leq i \leq k - 1, 0 \leq j \leq 4k^2 - 1\}$$

and

$$L = \left\{ y_{a,b} \middle| a, b \in \mathbb{Z}^2, 0 \leq a \leq 2k-1, 0 \leq b \leq 2k^2 - 1 \right\}.$$

Since $0 \leq i < k$ implies $b \leq ai + b < ak + b < 2k^2 + 2k^2 = 4k^2$ for all $0 \leq a < 2k$ and $0 \leq b < 2k^2$, the line $l_{a,b}$ contains the k points $(0,b), (1, a+b), \ldots, (k, ak+b)$ in P. Therefore $I(n,n) \geq I(P,L) = k|L| = kn = \frac{1}{\sqrt[3]{4}} n^{4/3}$.

For general n, we can use that $n \mapsto I(n,n)$ is monotone increasing, so

$$I(n,n) \geq I\left(4\left\lfloor \sqrt[3]{n/4} \right\rfloor^3\right) \geq \frac{1}{\sqrt[3]{4}}\left(4\left\lfloor \sqrt[3]{n/4} \right\rfloor^3\right)^{4/3} = 4\left\lfloor \sqrt[3]{n/4} \right\rfloor^4,$$

therefore

$$\liminf_{n\to\infty} \frac{I(n,n)}{n^{4/3}} \geq \liminf_{n\to\infty} \frac{4\left\lfloor \sqrt[3]{n/4} \right\rfloor^4}{n^{4/3}} \geq \liminf_{n\to\infty} \frac{4\left(\sqrt[3]{n/4} - 1\right)^4}{n^{4/3}} = \frac{1}{\sqrt[3]{4}}.$$

□

Recall the definition of a drawing of a graph (Definition 2.46).

Definition 3.4. Let d, γ be a drawing of $G = (V,E)$. For $x \in \mathbb{R}^2$, let $k_x = |\{e \in E | x \in \gamma_e((0,1))\}|$, i.e., the number of edges incident with x. The *crossing number* of the drawing is $\sum_{x\in\mathbb{R}^2} \binom{k_x}{2}$. The *crossing number* of a graph G, denoted by $\operatorname{cr}(G)$, is the minimum of the crossing numbers over its drawings.

Note that the endpoints of the curve γ_e are excluded, therefore $k_x = 0$ when $x = d(v)$ for a vertex $v \in v$.

Planar graphs are exactly the graphs satisfying $\operatorname{cr}(G) = 0$. It is known that a (simple) planar graph satifies $|E| \leq 3|V| - 6$. This can be improved to a lower bound on $\operatorname{cr}(G)$ in terms of $|V|$ and $|E|$.

Lemma 3.5. *Let $G = (V,E)$ be a simple graph. Then $\operatorname{cr}(G) \geq |E| - 3|V|$.*

Proof. Consider a drawing of G with $\operatorname{cr}(G)$ crossings, and delete one edge at each crossing. Then we obtain a planar graph with $|V|$ vertices and $|E| - \operatorname{cr}(G)$ edges, therefore $|E| - \operatorname{cr}(G) \leq 3|V| - 6$, which implies the claim. □

The following bound improves on Lemma 3.5 when $|E| > 8|V|$, and is asymptotically tight.

Theorem 3.6 (Crossing number theorem)**.** *Let $G = (V,E)$ be a simple graph. Then*

$$\operatorname{cr}(G) \geq \frac{1}{64}\frac{|E|^3}{|V|^2} - |V|.$$

Proof. If $|E| \leq 4|V|$, then the right hand side is at most 0, therefore the inequality is true. Assume now that $|E| > 4|V|$. Let $p = \frac{4|V|}{|E|} \in (0,1)$, and let $V' \subset V$ be a random subset formed by including each vertex of V with probability p, independently of each other. Let $G' = (V', E')$ be the (random) subgraph of G induced on V'.

Let (d, γ) be a drawing of G with crossing number $\operatorname{cr}(G)$. We may assume that there are no crossing edges incident with the same vertex. By restricting d and γ, we obtain a drawing of G' with some crossing number x'. By Lemma 3.5, we have

$$x' \geq \operatorname{cr}(G') \geq |E'| - 3|V'|,$$

therefore also $\mathbb{E}(x') \geq \mathbb{E}|E'| - 3\,\mathbb{E}|V'|$.

Each vertex of V appears in V' with probability p, an edge is present in E' if and only if both endpoints are in V', and a crossing of G gives rise to a crossing of G' if and only if both edges are in E'. It follows that

$$\begin{aligned}
\mathbb{E}|V'| &= p|V| \\
\mathbb{E}|E'| &= p^2|E| \\
\mathbb{E}(x') &= p^4 \operatorname{cr}(G).
\end{aligned}$$

By the previous estimates and using the choice of p, we obtain

$$\begin{aligned}
\operatorname{cr}(G) &= p^{-4}\,\mathbb{E}(x') \\
&\geq p^{-4}\,\mathbb{E}|E'| - 3p^{-4}\,\mathbb{E}|V'| \\
&\geq p^{-2}|E| - 3p^{-3}|V| \\
&= \frac{|E|^2}{16|V|^2}|E| - 3\frac{|E|^3}{64|V|^3}|V| \\
&= \frac{1}{64}\frac{|E|^3}{|V|^2} \\
&\geq \frac{1}{64}\frac{|E|^3}{|V|^2} - |V|
\end{aligned}$$

□

Theorem 3.7 (Szemerédi–Trotter)**.** *The maximum number of incidences between m points and n lines in the plane satisfies $I(m,n) \leq \sqrt[3]{32}m^{2/3}n^{2/3} + 4m + n$. In particular, $I(m,n) = O(m^{2/3}n^{2/3} + m + n)$.*

Proof. Let P be a set of points and L a set of lines in $\mathbb{R}^2$. We construct a graph $G = (V, E)$ as follows: the vertex set is $V = P$, and $\{u, v\}$ is an edge if both are incident with a line in L and $P \cap (u, v) = \emptyset$ (i.e., there are no points on the line between the two). The points and line segments then determine a drawing of G.

A line $l \in L$ containing k points corresponds to $k-1$ edges of the graph, therefore $I(P, L) = |E| + |L|$. The edges are subsets of $|L|$ lines, and the lines have at most $\binom{|L|}{2}$

intersection points, therefore $\operatorname{cr}(G) \le \binom{|L|}{2}$. From the crossing number theorem we have

$$\frac{1}{64}\frac{|E|^3}{|V|^2} - |V| \le \operatorname{cr}(G) \le \binom{|L|}{2},$$

which implies $|E|^3 \le 32|V|^2|L|^2 + 64|V|^3$. Since $a^3 + b^3 \le (a+b)^3$, we obtain

$$I(P,L) = |E| + |L| \le \sqrt[3]{32}|P|^{2/3}|L|^{2/3} + 4|P| + |L|.$$

□

Exercise 3.1. Prove that, for all $m, n \in \mathbb{N}$, the inequality $I(m,n) \ge m + n - 1$ holds.

3.2 Unit distances

We study distances in point sets.

Definition 3.8. Let P be a set of points in the plane. The number of *unit distances* is $U(P) = \left|\left\{\{x,y\} \in \binom{P}{2} \middle| \operatorname{dist}(x,y) = 1\right\}\right|$, and the number of *distinct distances* if $g(P) = |\{\operatorname{dist}(x,y) | x, y \in P, x \ne y\}|$.

For $n \in \mathbb{N}$, we denote by $U(n)$ the maximum of $U(P)$ over n-point sets P, and by $g(n)$ the minimum of $g(P)$ over n-point sets P.

Clearly, $U(n) \le \binom{n}{2}$. For small values of n, one can verify that $U(1) = 0$, $U(2) = 1$, $U(3) = 3$, $U(4) = 5$. The two quantities are related as $U(n)g(n) \ge \binom{n}{2}$.

An upper bound on $U(n)$ can be obtained in a similar way as on $I(n,n)$, which gives $U(n) = O(n^{4/3})$, and as a consequence $g(n) = \Omega(n^{2/3})$. The precise order of magnitude is not known. We explain the best known lower bound, starting with number theoretic tools.

The configuration used in the lower bound consists of the vertices of an (appropriately rescaled) square grid. To understand the number of occurrences of a given distance, we look for integers that can be written as a sum of squares in many ways.

Lemma 3.9. *Let $p_1, p_2, \ldots, p_r$ be distinct primes of the form $4k+1$ ($k \in \mathbb{N}$). Then $N = p_1 p_2 \ldots p_r$ can be written as a sum of squares in at least 2^r ways.*

Proof. We work in the ring of Gaussian integers $\mathbb{Z}_i$, which is a Euclidean domain, therefore has unique factorization: every Gaussian integer is the product of a unit and Gaussian primes, in a unique way up to reordering and multiplying the factors by units (associates). The units are $\pm 1, \pm i$, and the primes are $1+i$, the primes in $\mathbb{Z}$ of the form $4k+3$, and exactly two factors (which are necessarily conjugates of each other) of the primes in $\mathbb{Z}$ of the form $4k+1$.

Write $p_j = \tau_j \overline{\tau_j}$ with $\tau_j \in \mathbb{Z}_j$. For every subset $J \subseteq [r]$ we have a factorization

$$N = \underbrace{\left(\prod_{j\in J} \tau_j \prod_{j \in [r]\setminus J} \overline{\tau_j}\right)}_{A_J + iB_j} \underbrace{\left(\prod_{j\in J} \overline{\tau_j} \prod_{j \in [r]\setminus J} \tau_j\right)}_{A_J - iB_J},$$

which also means $N = A_J^2 + B_J^2$. By the unique factorization property, the factor $A_J + iB_j$ determines the subset J as well. The only ambiguity is the order of the two factors, therefore the 2^r subsets determine 2^{r-1} decompositions. Accounting for the possibility of swapping the two terms, we get 2^r sums if the order matters. □

Two integers a and b are *coprime* if their only positive common divisor is 1. Recall that *Euler's function* φ is the number theoretic function defined such that $\varphi(d)$ is the number of integers in $[1, d]$ that are coprime to d.

Theorem 3.10. *Let $a, d \in \mathbb{N}_{>0}$ be coprime and let $\pi_{d,a}(n)$ denote the number of primes in $[1, n]$ which are of the form $a + kd$ for some $k \in \mathbb{N}$. Then*

$$\pi_{d,a}(n) = (1 + o(1)) \frac{1}{\varphi(d)} \frac{n}{\ln n}.$$

The prime number theorem is the special case $a = d = 1$.

Theorem 3.11. *There is a constant c such that for all $n \geq 2$, $U(n) \geq n^{1+\frac{c}{\log \log n}}$.*

Proof. Let p_j denote the jth prime of the form $4k + 1$, and let r be the largest number such that $M = p_1 p_2 \cdots p_{r-1} \leq \frac{n}{4}$. Suppose that n is a square, and let P be the points of a $\sqrt{n} \times \sqrt{n}$ grid with step size $\frac{1}{\sqrt{M}}$.

For each point $p \in P$, there at least as many points $q \in P$ with $\text{dist}(p, q) = 1$ as the number of (ordered) ways to write M as a sum of two squares of positive integers. Since there are 4 sign choices, there are at least $\frac{2^{r-1}}{4}$ such points q, and at least $\frac{n2^{r-1}}{8}$ (unordered) pairs of points at distance 1 in total.

Using the choice of r and $5 = p_1 \leq p_j \leq p_r$, we obtain the chain of inequalities

$$e^r \leq 4p_1 p_2 \cdots p_{r-1} \leq n < 4p_1 p_2 \cdots p_{r-1} p_r < 4p_r^r.$$

Theorem 3.10 implies that

$$r = \pi_{4,1}(p_r) = (1 + o(1)) \frac{1}{2} \frac{p_r}{\ln p_r} \geq \sqrt{p_r} \geq n^{1/3r}$$

for large n. This implies

$$r \geq \frac{\ln n}{3 \log r} \geq \frac{\ln n}{3 \ln \ln n}$$

and therefore

$$U(n) \geq U(P) \geq n2^{r-4} \geq n^{1+\frac{c}{\ln \ln n}}$$

when n is large. □