Introduction to Discrete Geometry

Lecture Notes

Jiří Matoušek
Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, Czech Republic

September 2003

1

Convexity

We begin with a review of basic geometric notions such as hyperplanes and affine subspaces in $\mathbb{R}^d$, and we spend some time by discussing the notion of general position. Then we consider fundamental properties of convex sets in $\mathbb{R}^d$, such as a theorem about the separation of disjoint convex sets by a hyperplane and Helly's theorem.

1.1 Linear and Affine Subspaces, General Position

Linear subspaces. Let $\mathbb{R}^d$ denote the d-dimensional Euclidean space. The points are d-tuples of real numbers, $x = (x_1, x_2, \ldots, x_d)$.

The space $\mathbb{R}^d$ is a vector space, and so we may speak of linear subspaces, linear dependence of points, linear span of a set, and so on. A linear subspace of $\mathbb{R}^d$ is a subset closed under addition of vectors and under multiplication by real numbers. What is the geometric meaning? For instance, the linear subspaces of $\mathbb{R}^2$ are the origin itself, all lines passing through the origin, and the whole of $\mathbb{R}^2$. In $\mathbb{R}^3$, we have the origin, all lines and planes passing through the origin, and $\mathbb{R}^3$.

Affine notions. An arbitrary line in $\mathbb{R}^2$, say, is *not* a linear subspace unless it passes through 0. General lines are what are called *affine subspaces*. An affine subspace of $\mathbb{R}^d$ has the form $x+L$, where $x \in \mathbb{R}^d$ is some

vector and L is a linear subspace of $\mathbb{R}^d$. Having defined affine subspaces, the other "affine" notions can be constructed by imitating the "linear" notions.

What is the *affine hull* of a set $X \subseteq \mathbb{R}^d$? It is the intersection of all affine subspaces of $\mathbb{R}^d$ containing X. As is well known, the linear span of a set X can be described as the set of all linear combinations of points of X. What is an *affine combination* of points $a_1, a_2, \ldots, a_n \in \mathbb{R}^d$ that would play an analogous role? To see this, we translate the whole set by $-a_n$, so that a_n becomes the origin, we make a linear combination, and we translate back by $+a_n$. This yields an expression of the form $\beta_1(a_1 - a_n) + \beta_2(a_2 - a_n) + \cdots + \beta_n(a_n - a_n) + a_n = \beta_1 a_1 + \beta_2 a_2 + \cdots + \beta_{n-1}a_{n-1} + (1 - \beta_1 - \beta_2 - \cdots - \beta_{n-1})a_n$, where $\beta_1, \ldots, \beta_n$ are arbitrary real numbers. Thus, an affine combination of points $a_1, \ldots, a_n \in \mathbb{R}^d$ is an expression of the form

$$\alpha_1 a_1 + \cdots + \alpha_n a_n, \text{ where } \alpha_1, \ldots, \alpha_n \in \mathbb{R} \text{ and } \alpha_1 + \cdots + \alpha_n = 1.$$

Then indeed, it is not hard to check that the affine hull of X is the set of all affine combinations of points of X.

The *affine dependence* of points $a_1, \ldots, a_n$ means that one of them can be written as an affine combination of the others. This is the same as the existence of real numbers $\alpha_1, \alpha_2, \ldots \alpha_n$, at least one of them nonzero, such that both

$$\alpha_1 a_1 + \alpha_2 a_2 + \cdots + \alpha_n a_n = 0 \text{ and } \alpha_1 + \alpha_2 + \cdots + \alpha_n = 0.$$

(Note the difference: In an affine *combination*, the α_i sum to 1, while in an affine *dependence*, they sum to 0.)

Affine dependence of $a_1, \ldots, a_n$ is equivalent to linear dependence of the $n-1$ vectors $a_1 - a_n, a_2 - a_n, \ldots, a_{n-1} - a_n$. Therefore, the maximum possible number of affinely independent points in $\mathbb{R}^d$ is $d+1$.

Another way of expressing affine dependence uses "lifting" one dimension higher. Let $b_i = (a_i, 1)$ be the vector in $\mathbb{R}^{d+1}$ obtained by appending a new coordinate equal to 1 to a_i; then $a_1, \ldots, a_n$ are affinely dependent if and only if $b_1, \ldots, b_n$ are linearly dependent. This correspondence of

affine notions in $\mathbb{R}^d$ with linear notions in $\mathbb{R}^{d+1}$ is quite general. For example, if we identify $\mathbb{R}^2$ with the plane $x_3 = 1$ in $\mathbb{R}^3$ as in the picture,

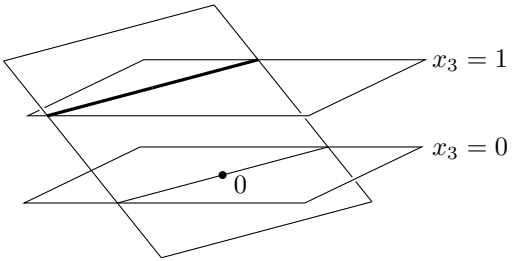

then we obtain a bijective correspondence of the k-dimensional linear subspaces of $\mathbb{R}^3$ that do not lie in the plane $x_3 = 0$ with $(k-1)$-dimensional affine subspaces of $\mathbb{R}^2$. The drawing shows a 2-dimensional linear subspace of $\mathbb{R}^3$ and the corresponding line in the plane $x_3 = 1$. (The same works for affine subspaces of $\mathbb{R}^d$ and linear subspaces of $\mathbb{R}^{d+1}$ not contained in the subspace $x_{d+1} = 0$.)

Let $a_1, a_2, \ldots, a_{d+1}$ be points in $\mathbb{R}^d$, and let A be the $d \times d$ matrix with $a_i - a_{d+1}$ as the ith column, $i = 1, 2, \ldots, d$. Then $a_1, \ldots, a_{d+1}$ are affinely independent if and only if A has d linearly independent columns, and this is equivalent to $\det(A) \neq 0$. We have a useful criterion of affine independence using a determinant.

Affine subspaces of $\mathbb{R}^d$ of certain dimensions have special names. A $(d-1)$-dimensional affine subspace of $\mathbb{R}^d$ is called a *hyperplane* (while the word *plane* usually means a 2-dimensional subspace of $\mathbb{R}^d$ for any d). One-dimensional subspaces are lines, and a k-dimensional affine subspace is often called a *k-flat*.

A hyperplane is usually specified by a single linear equation of the form $a_1x_1 + a_2x_2 + \cdots + a_dx_d = b$. We usually write the left-hand side as the scalar product $\langle a, x\rangle$. So a hyperplane can be expressed as the set $\{x \in \mathbb{R}^d\colon \langle a, x\rangle = b\}$ where $a \in \mathbb{R}^d \setminus \{0\}$ and $b \in \mathbb{R}$. A (closed) *half-space* in $\mathbb{R}^d$ is a set of the form $\{x \in \mathbb{R}^d\colon \langle a, x\rangle \geq b\}$ for some $a \in \mathbb{R}^d \setminus \{0\}$; the hyperplane $\{x \in \mathbb{R}^d\colon \langle a, x\rangle = b\}$ is its boundary.

General k-flats can be given either as intersections of hyperplanes or as affine images of $\mathbb{R}^k$ (parametric expression). In the first case, an

intersection of k hyperplanes can also be viewed as a solution to a system $Ax = b$ of linear equations, where $x \in \mathbb{R}^d$ is regarded as a column vector, A is a $k \times d$ matrix, and $b \in \mathbb{R}^k$. (As a rule, in formulas involving matrices, we interpret points of $\mathbb{R}^d$ as *column* vectors.)

An *affine mapping* $f\colon \mathbb{R}^k \to \mathbb{R}^d$ has the form $f\colon y \mapsto By + c$ for some $d \times k$ matrix B and some $c \in \mathbb{R}^d$, so it is a composition of a linear map with a translation. The image of f is a k'-flat for some $k' \le \min(k, d)$. This k' equals the rank of the matrix B.

General position. *"We assume that the points (lines, hyperplanes,...) are in general position."* This magical phrase appears in many proofs. Intuitively, general position means that no "unlikely coincidences" happen in the considered configuration. For example, if 3 points are chosen in the plane without any special intention, "randomly," they are unlikely to lie on a common line. For a planar point set in general position, we always require that no three of its points be collinear. For points in $\mathbb{R}^d$ in general position, we assume similarly that no unnecessary affine dependencies exist: No $k \le d+1$ points lie in a common $(k-2)$-flat. For lines in the plane in general position, we postulate that no 3 lines have a common point and no 2 are parallel.

The precise meaning of general position is not fully standard: It may depend on the particular context, and to the usual conditions mentioned above we sometimes add others where convenient. For example, for a planar point set in general position we can also suppose that no two points have the same x-coordinate.

What conditions are suitable for including into a "general position" assumption? In other words, what can be considered as an unlikely coincidence? For example, let X be an n-point set in the plane, and let the coordinates of the ith point be (x_i, y_i). Then the vector $v(X) = (x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n)$ can be regarded as a point of $\mathbb{R}^{2n}$. For a configuration X in which $x_1 = x_2$, i.e., the first and second points have the same x-coordinate, the point $v(X)$ lies on the hyperplane $\{x_1 = x_2\}$ in $\mathbb{R}^{2n}$. The configurations X where *some* two points share the x-coordinate thus correspond to the union of $\binom{n}{2}$ hyperplanes in $\mathbb{R}^{2n}$. Since a hyperplane in $\mathbb{R}^{2n}$ has ($2n$-dimensional) measure zero, almost all points of $\mathbb{R}^{2n}$

correspond to planar configurations X with all the points having distinct x-coordinates. In particular, if X is any n-point planar configuration and $\varepsilon > 0$ is any given real number, then there is a configuration X', obtained from X by moving each point by distance at most ε, such that all points of X' have distinct x-coordinates. Not only that: almost all small movements (*perturbations*) of X result in X' with this property.

This is the key property of general position: configurations in general position lie arbitrarily close to any given configuration (and they abound in any small neighborhood of any given configuration). Here is a fairly general type of condition with this property. Suppose that a configuration X is specified by a vector $t = (t_1, t_2, \ldots, t_m)$ of m real numbers (coordinates). The objects of X can be points in $\mathbb{R}^d$, in which case $m = dn$ and the t_j are the coordinates of the points, but they can also be circles in the plane, with $m = 3n$ and the t_j expressing the center and the radius of each circle, and so on. The general position condition we can put on the configuration X is $p(t) = p(t_1, t_2, \ldots, t_m) \neq 0$, where p is some nonzero polynomial in m variables. Here we use the following well-known fact: *For any nonzero m-variate polynomial $p(t_1, \ldots, t_m)$, the zero set $\{t \in \mathbb{R}^m\colon p(t) = 0\}$ has measure 0 in $\mathbb{R}^m$.*

Therefore, almost all configurations X satisfy $p(t) \neq 0$. So any condition that can be expressed as $p(t) \neq 0$ for a certain polynomial p in m real variables, or, more generally, as $p_1(t) \neq 0$ or $p_2(t) \neq 0$ or $\ldots$, for finitely or countably many polynomials $p_1, p_2, \ldots$, can be included in a general position assumption.

For example, let X be an n-point set in $\mathbb{R}^d$, and let us consider the condition "no $d{+}1$ points of X lie in a common hyperplane." In other words, no $d{+}1$ points should be affinely dependent. As we know, the affine dependence of $d{+}1$ points means that a suitable $d \times d$ determinant equals 0. This determinant is a polynomial (of degree d) in the coordinates of these $d{+}1$ points. Introducing one polynomial for every $(d{+}1)$-tuple of the points, we obtain $\binom{n}{d+1}$ polynomials such that at least one of them is 0 for any configuration X with $d{+}1$ points in a common hyperplane. Other usual conditions for general position can be expressed similarly.

In many proofs, assuming general position simplifies matters considerably. But what do we do with configurations X_0 that are not in general position? We have to argue, somehow, that if the statement being proved is valid for configurations X arbitrarily close to our X_0, then it must be valid for X_0 itself, too. Such proofs, usually called *perturbation arguments*, are often rather simple, and almost always somewhat boring. But sometimes they can be tricky, and one should not underestimate them, no matter how tempting this may be. A nontrivial example will be demonstrated in Section 5.5 (Lemma 5.5.4).

1.2 Convex Sets, Convex Combinations, Separation

Intuitively, a set is convex if its surface has no "dips":

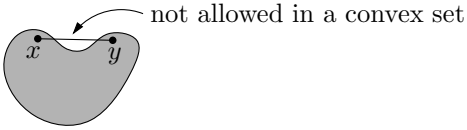

1.2.1 Definition (Convex set). *A set $C \subseteq \mathbb{R}^d$ is* convex *if for every two points $x, y \in C$ the whole segment xy is also contained in C. In other words, for every $t \in [0,1]$, the point $tx + (1-t)y$ belongs to C.*

The intersection of an arbitrary family of convex sets is obviously convex. So we can define the *convex hull* of a set $X \subseteq \mathbb{R}^d$, denoted by $\mathrm{conv}(X)$, as the intersection of all convex sets in $\mathbb{R}^d$ containing X. Here is a planar example with a finite X:

An alternative description of the convex hull can be given using convex combinations.

1.2.2 Claim. *A point x belongs to* $\mathrm{conv}(X)$ *if and only if there exist points $x_1, x_2, \dots x_n \in X$ and nonnegative real numbers $t_1, t_2, \dots, t_n$ with $\sum_{i=1}^n t_i = 1$ such that $x = \sum_{i=1}^n t_i x_i$.*

The expression $\sum_{i=1}^n t_i x_i$ as in the claim is called a *convex combination* of the points $x_1, x_2, \dots, x_n$. (Compare this with the definitions of linear and affine combinations.)

Sketch of proof. Each convex combination of points of X must lie in $\mathrm{conv}(X)$: For $n = 2$ this is by definition, and for larger n by induction. Conversely, the set of all convex combinations obviously contains X and it is convex. □

In $\mathbb{R}^d$, it is sufficient to consider convex combinations involving at most $d+1$ points:

1.2.3 Theorem (Carathéodory's theorem). *Let $X \subseteq \mathbb{R}^d$. Then each point of* $\mathrm{conv}(X)$ *is a convex combination of at most $d+1$ points of X.*

For example, in the plane, $\mathrm{conv}(X)$ is the union of all triangles with vertices at points of X. The proof of the theorem is left as an exercise.

A basic result about convex sets is the separability of disjoint convex sets by a hyperplane.

1.2.4 Theorem (Separation theorem). *Let $C, D \subseteq \mathbb{R}^d$ be convex sets with $C \cap D = \emptyset$. Then there exists a hyperplane h such that C lies in one of the closed half-spaces determined by h, and D lies in the opposite closed half-space. In other words, there exist a unit vector $a \in \mathbb{R}^d$ and a number $b \in \mathbb{R}$ such that for all $x \in C$ we have $\langle a, x\rangle \geq b$, and for all $x \in D$ we have $\langle a, x\rangle \leq b$.*

If C and D are closed and at least one of them is bounded, they can be separated strictly; in such a way that $C \cap h = D \cap h = \emptyset$.

In particular, a closed convex set can be strictly separated from a point. This implies that the convex hull of a closed set X equals the intersection of all closed half-spaces containing X.

Sketch of proof. First assume that C and D are compact (i.e., closed and bounded). Then the Cartesian product $C \times D$ is a compact space, too, and the distance function $(x, y) \mapsto \|x - y\|$ attains its minimum on $C \times D$. That is, there exist points $p \in C$ and $q \in D$ such that the distance of C and D equals the distance of p and q.

The desired separating hyperplane h can be taken as the one perpendicular to the segment pq and passing through its midpoint:

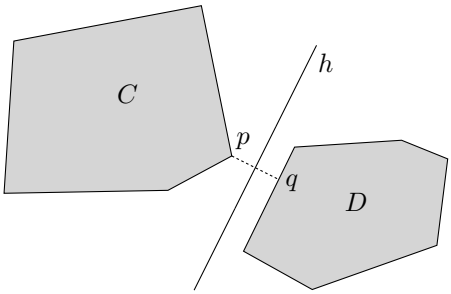

It is easy to check that h indeed avoids both C and D.

If D is compact and C closed, we can intersect C with a large ball and get a compact set C'. If the ball is sufficiently large, then C and C' have the same distance to D. So the distance of C and D is attained at some $p \in C'$ and $q \in D$, and we can use the previous argument.

For arbitrary disjoint convex sets C and D the proof takes more work. First we simplify the situation by observing that the hyperplane separation of C and D is equivalent to a hyperplane separation of the convex set $C - D = \{x - y\colon x \in C, y \in D\}$ from the point 0 (we leave the equivalence as an easy exercise; we could have used this trick in the first part of the proof as well). Hence it suffices to consider the case with $D = \{0\}$ and C an arbitrary convex set not containing 0.

If $0 \notin \overline{C}$, where $\overline{C}$ denotes the closure of C, then we are done by the first part of the proof, so we assume $0 \in \overline{C}$. We claim that 0 must be a boundary point of $\overline{C}$. The proof of this proceeds by contradiction:If 0 is not on the boundary, then there is an open ball B around 0 with $B \subseteq \overline{C}$. By scaling the coordinate system suitably, we may assume that B contains the points $p_1, p_2, \ldots, p_{d+1}$,where $p_i = e_i$ (the ith standard

unit vector) for $i = 1, 2, \ldots, d$ and $p_{d+1} = -(e_1 + e_2 + \cdots + e_d)$. It is very intuitive, and a simple calculation confirms it, that $\operatorname{conv}\{p_1, \ldots, p_{d+1}\}$ contains a small neighborhood of 0. Since $B \subseteq \overline{C}$, for every $\varepsilon > 0$ there is a point $p_i' \in C$ lying at distance at most ε from p_i, $i = 1, \ldots, d+1$, and it is again routine to check that if $\varepsilon > 0$ is sufficiently small, $\operatorname{conv}\{p_1', \ldots, p_{d+1}'\} \subseteq C$ contains 0. But this is a contradiction since we assumed $0 \notin C$. The claim is verified.

We now know that 0 lies on the boundary of $\overline{C}$, and hence in the closure of the complement of $\overline{C}$. We can thus choose a sequence $\{q_n\}_{n=1}^{\infty}$ of points in $\mathbb{R}^d \setminus \overline{C}$ tending to 0.Each q_n lies outside the closed convex set $\overline{C}$, and thus it can be separated from it (strictly) by a hyperplane $h_n = \{x \in \mathbb{R}^d \colon \langle a_n, x\rangle = b_n\}$, where a_n is a unit vector and $b_n \in \mathbb{R}$. The sequence $(b_n)_{n=1}^{\infty}$ is bounded (this also needs a small argument), and by compactness, the sequence of $(d+1)$-component vectors $(a_n, b_n) \in \mathbb{R}^{d+1}$ has a cluster point (a, b). One can verify, by contradiction, that the hyperplane $h = \{x \in \mathbb{R}^d \colon \langle a, x\rangle = b\}$ separates the point 0 from C (non-strictly). □

The importance of the separation theorem is documented by its presence in several branches of mathematics in various disguises. Its home territory is probably functional analysis, where it is formulated and proved for infinite-dimensional spaces; essentially it is the so-called Hahn–Banach theorem. The usual functional-analytic proof is different from the one we gave, and in a way it is more elegant and conceptual. The proof sketched above uses more special properties of $\mathbb{R}^d$, but it is quite short and intuitive in the case of compact C and D.

Connection to linear programming. A basic result in the theory of linear programming is the Farkas lemma. It is a special case of the duality of linear programming as well as the key step in its proof.

1.2.5 Lemma (Farkas lemma, one of many versions). *For every $d \times n$ real matrix A, exactly one of the following cases occurs:*

(i) *The system of linear equations $Ax = 0$ has a nontrivial nonnegative solution $x \in \mathbb{R}^n$ (all components of x are nonnegative and at least*

one of them is strictly positive).

(ii) *There exists a $y \in \mathbb{R}^d$ such that $y^T A$ is a vector with all entries strictly negative. Thus, if we multiply the jth equation in the system $Ax = 0$ by y_j and add these equations together, we obtain an equation that* obviously *has no nontrivial nonnegative solution, since all the coefficients on the left-hand sides are strictly negative, while the right-hand side is* 0.

Proof. Let us see why this is yet another version of the separation theorem. Let $V \subset \mathbb{R}^d$ be the set of n points given by the column vectors of the matrix A. We distinguish two cases: Either $0 \in \operatorname{conv}(V)$ or $0 \notin \operatorname{conv}(V)$.

In the former case, we know that 0 is a convex combination of the points of V, and the coefficients of this convex combination determine a nontrivial nonnegative solution to $Ax = 0$.

In the latter case, there exists a hyperplane strictly separating V from 0, i.e., a unit vector $y \in \mathbb{R}^d$ such that $\langle y, v\rangle < \langle y, 0\rangle = 0$ for each $v \in V$. This is just the y from the second alternative in the Farkas lemma. □

1.3 Radon's Lemma and Helly's Theorem

Carathéodory's theorem from the previous section, together with Radon's lemma and Helly's theorem presented here, are three basic properties of convexity in $\mathbb{R}^d$ involving the dimension. We begin with Radon's lemma.

1.3.1 Theorem (Radon's lemma). *Let A be a set of $d+2$ points in $\mathbb{R}^d$. Then there exist two* disjoint *subsets $A_1, A_2 \subset A$ such that*

$$\operatorname{conv}(A_1) \cap \operatorname{conv}(A_2) \neq \emptyset.$$

A point $x \in \operatorname{conv}(A_1) \cap \operatorname{conv}(A_2)$, where A_1 and A_2 are as in the theorem, is called a *Radon point* of A, and the pair (A_1, A_2) is called a *Radon partition* of A (it is easily seen that we can require $A_1 \cup A_2 = A$).

Here are two possible cases in the plane:

Proof. Let $A = \{a_1, a_2, \ldots, a_{d+2}\}$. These $d+2$ points are necessarily affinely dependent. That is, there exist real numbers $\alpha_1, \ldots, \alpha_{d+2}$, not all of them 0, such that $\sum_{i=1}^{d+2} \alpha_i = 0$ and $\sum_{i=1}^{d+2} \alpha_i a_i = 0$.

Set $P = \{i\colon \alpha_i > 0\}$ and $N = \{i\colon \alpha_i < 0\}$. Both P and N are nonempty. We claim that P and N determine the desired subsets. Let us put $A_1 = \{a_i\colon i \in P\}$ and $A_2 = \{a_i\colon i \in N\}$. We are going to exhibit a point x that is contained in the convex hulls of both these sets.

Put $S = \sum_{i\in P} \alpha_i$; we also have $S = -\sum_{i\in N} \alpha_i$. Then we define

$$x = \sum_{i\in P} \frac{\alpha_i}{S}\, a_i\,. \tag{1.1}$$

Since $\sum_{i=1}^{d+2} \alpha_i a_i = 0 = \sum_{i\in P} \alpha_i a_i + \sum_{i\in N} \alpha_i a_i$, we also have

$$x = \sum_{i\in N} \frac{-\alpha_i}{S}\, a_i\,. \tag{1.2}$$

The coefficients of the a_i in (1.1) are nonnegative and sum to 1, so x is a convex combination of points of A_1. Similarly (1.2) expresses x as a convex combination of points of A_2. □

Helly's theorem is one of the most famous results of a combinatorial nature about convex sets.

1.3.2 Theorem (Helly's theorem). *Let $C_1, C_2, \ldots, C_n$ be convex sets in $\mathbb{R}^d$, $n \geq d+1$. Suppose that the intersection of every $d+1$ of these sets is nonempty. Then the intersection of all the C_i is nonempty.*

The first nontrivial case states that if every 3 among 4 convex sets in the plane intersect, then there is a point common to all 4 sets. This can be proved by an elementary geometric argument, perhaps distinguishing a few cases, and the reader may want to try to find a proof before reading further.

In a contrapositive form, Helly's theorem guarantees that whenever $C_1, C_2, \ldots, C_n$ are convex sets with $\bigcap_{i=1}^{n} C_i = \emptyset$, then this is witnessed by some at most $d+1$ sets with empty intersection among the C_i. In this way, many proofs are greatly simplified, since in planar problems, say, one can deal with 3 convex sets instead of an arbitrary number.

It is very tempting and quite usual to formulate Helly's theorem as follows: "If every $d+1$ among n convex sets in $\mathbb{R}^d$ intersect, then all the sets intersect." But, strictly speaking, this is false, for a trivial reason: For $d \geq 2$, the assumption as stated here is met by $n = 2$ disjoint convex sets.

Proof of Helly's theorem. (Using Radon's lemma.) For a fixed d, we proceed by induction on n. The case $n = d+1$ is clear, so we suppose that $n \geq d+2$ and that the statement of Helly's theorem holds for smaller n. Actually, $n = d+2$ is the crucial case; the result for larger n follows at once by a simple induction.

Consider sets $C_1, C_2, \ldots, C_n$ satisfying the assumptions. If we leave out any one of these sets, the remaining sets have a nonempty intersection by the inductive assumption. Let us fix a point $a_i \in \bigcap_{j \neq i} C_j$ and consider the points $a_1, a_2, \ldots, a_{d+2}$. By Radon's lemma, there exist disjoint index sets $I_1, I_2 \subset \{1, 2, \ldots, d+2\}$ such that

$$\operatorname{conv}(\{a_i\colon i \in I_1\}) \cap \operatorname{conv}(\{a_i\colon i \in I_2\}) \neq \emptyset.$$

We pick a point x in this intersection. The following picture illustrates the case $d = 2$ and $n = 4$:

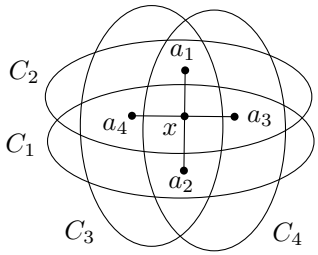

We claim that x lies in the intersection of all the C_i. Consider some $i \in \{1, 2, \ldots, n\}$; then $i \notin I_1$ or $i \notin I_2$. In the former case, each a_j

with $j \in I_1$ lies in C_i, and so $x \in \operatorname{conv}(\{a_j\colon j \in I_1\}) \subseteq C_i$. For $i \notin I_2$ we similarly conclude that $x \in \operatorname{conv}(\{a_j\colon j \in I_2\}) \subseteq C_i$. Therefore $x \in \bigcap_{i=1}^{n} C_i$. □

An infinite version of Helly's theorem. If we have an infinite collection of convex sets in $\mathbb{R}^d$ such that any $d+1$ of them have a common point, the entire collection still need not have a common point. Two examples in $\mathbb{R}^1$ are the families of intervals $\{(0, 1/n)\colon n = 1, 2, \ldots\}$ and $\{[n, \infty)\colon n = 1, 2, \ldots\}$. The sets in the first example are not closed, and the second example uses unbounded sets. For *compact* (i.e., closed and bounded) sets, the theorem holds:

1.3.3 Theorem (Infinite version of Helly's theorem). *Let $\mathcal{C}$ be an arbitrary infinite family of compact convex sets in $\mathbb{R}^d$ such that any $d+1$ of the sets have a nonempty intersection. Then all the sets of $\mathcal{C}$ have a nonempty intersection.*

Proof. By Helly's theorem, any finite subfamily of $\mathcal{C}$ has a nonempty intersection. By a basic property of compactness, if we have an arbitrary family of compact sets such that each of its finite subfamilies has a nonempty intersection, then the entire family has a nonempty intersection. □

1.4 Centerpoint and Ham Sandwich

We prove an interesting result as an application of Helly's theorem.

1.4.1 Definition (Centerpoint). *Let X be an n-point set in $\mathbb{R}^d$. A point $x \in \mathbb{R}^d$ is called a* centerpoint *of X if each closed half-space containing x contains at least $\frac{n}{d+1}$ points of X.*

Let us stress that one set may generally have many centerpoints, and a centerpoint need not belong to X.

The notion of centerpoint can be viewed as a generalization of the *median* of one-dimensional data. Suppose that $x_1, \dots, x_n \in \mathbb{R}$ are results of measurements of an unknown real parameter x. How do we estimate x from the x_i? We can use the arithmetic mean, but if one of the measurements is completely wrong (say, 100 times larger than the others), we may get quite a bad estimate. A more "robust" estimate is a *median*, i.e., a point x such that at least $\frac{n}{2}$ of the x_i lie in the interval $(-\infty, x]$ and at least $\frac{n}{2}$ of them lie in $[x, \infty)$. The centerpoint can be regarded as a generalization of the median for higher-dimensional data.

In the definition of centerpoint we could replace the fraction $\frac{1}{d+1}$ by some other parameter $\alpha \in (0, 1)$. For $\alpha > \frac{1}{d+1}$, such an "α-centerpoint" need not always exist: Take $d+1$ points in general position for X. With $\alpha = \frac{1}{d+1}$ as in the definition above, a centerpoint always exists, as we prove next.

Centerpoints are important, for example, in some algorithms of divide-and-conquer type, where they help divide the considered problem into smaller subproblems. Since no really efficient algorithms are known for finding "exact" centerpoints, the algorithms often use α-centerpoints with a suitable $\alpha < \frac{1}{d+1}$, which are easier to find.

1.4.2 Theorem (Centerpoint theorem). *Each finite point set in $\mathbb{R}^d$ has at least one centerpoint.*

Proof. First we note an *equivalent definition of a centerpoint:* x is a centerpoint of X if and only if it lies in each open half-space γ such that $|X \cap \gamma| > \frac{d}{d+1}\, n$.

We would like to apply Helly's theorem to conclude that all these open half-spaces intersect. But we cannot proceed directly, since we have infinitely many half-spaces and they are open and unbounded. Instead of such an open half-space γ, we thus consider the compact convex set $\operatorname{conv}(X \cap \gamma) \subset \gamma$.

Letting γ run through all open half-spaces γ with $|X \cap \gamma| > \frac{d}{d+1}\, n$, we obtain a family $\mathcal{C}$ of compact convex sets. Each of them contains more than $\frac{d}{d+1} n$ points of X, and so the intersection of any $d+1$ of them contains at least one point of X. The family $\mathcal{C}$ consists of finitely many distinct sets (since X has finitely many distinct subsets), and so $\bigcap \mathcal{C} \neq \emptyset$ by Helly's theorem. Each point in this intersection is a centerpoint. □

In the definition of a centerpoint we can regard the finite set X as defining a distribution of mass in $\mathbb{R}^d$. The centerpoint theorem asserts that for some point x, any half-space containing x encloses at least $\frac{1}{d+1}$ of the total mass. It is not difficult to show that this remains valid for continuous mass distributions, or even for arbitrary Borel probability measures on $\mathbb{R}^d$.

Ham-sandwich theorem and its relatives. Here is another important result, not much related to convexity but with a flavor resembling the centerpoint theorem.

1.4.3 Theorem (Ham-sandwich theorem). *Every d finite sets in $\mathbb{R}^d$ can be simultaneously bisected by a hyperplane. A hyperplane h* bisects *a finite set A if each of the open half-spaces defined by h contains at most $\lfloor |A|/2 \rfloor$ points of A.*

This theorem is usually proved via continuous mass distributions using a tool from algebraic topology: the *Borsuk–Ulam theorem*. Here we omit a proof.

Note that if A_i has an odd number of points, then every h bisecting A_i passes through a point of A_i. Thus if $A_1, \ldots, A_d$ all have odd sizes and their union is in general position, then every hyperplane simultaneously bisecting them is determined by d points, one of each A_i. In particular, there are only finitely many such hyperplanes.

Again, an analogous ham-sandwich theorem holds for arbitrary d Borel probability measures in $\mathbb{R}^d$.

Center transversal theorem. There can be beautiful new things to discover even in well-studied areas of mathematics. A good example is the following recent result, which "interpolates" between the centerpoint theorem and the ham-sandwich theorem.

1.4.4 Theorem (Center transversal theorem). *Let $1 \le k \le d$ and let $A_1, A_2, \ldots, A_k$ be finite point sets in $\mathbb{R}^d$. Then there exists a $(k-1)$-flat f such that for every hyperplane h containing f, both the closed half-spaces defined by h contain at least $\frac{1}{d-k+2}|A_i|$ points of A_i, $i = 1, 2, \ldots, k$.*

The ham-sandwich theorem is obtained for $k = d$ and the centerpoint theorem for $k = 1$. The proof, which we again have to omit, is based on a result of algebraic topology, too, but it uses a considerably more advanced machinery than the ham-sandwich theorem.

2

Lattices and Minkowski's Theorem

This chapter is a quick excursion into the *geometry of numbers*, a field where number-theoretic results are proved by geometric arguments, often using properties of convex bodies in $\mathbb{R}^d$. We formulate the simple but beautiful theorem of Minkowski on the existence of a nonzero lattice point in every symmetric convex body of sufficiently large volume. We derive several consequences, concluding with a geometric proof of the famous theorem of Lagrange claiming that every natural number can be written as the sum of at most 4 squares.

2.1 Minkowski's Theorem

In this section we consider the integer lattice $\mathbb{Z}^d$, and so a lattice point is a point in $\mathbb{R}^d$ with integer coordinates. The following theorem can be used in many interesting situations to establish the existence of lattice points with certain properties.

2.1.1 Theorem (Minkowski's theorem). *Let $C \subseteq \mathbb{R}^d$ be symmetric (around the origin, i.e., $C = -C$), convex, bounded, and suppose that $\mathrm{vol}(C) > 2^d$. Then C contains at least one lattice point different from 0.*

Proof. We put $C' = \frac{1}{2}C = \{\frac{1}{2}x\colon x \in C\}$.

Claim: There exists a nonzero integer vector $v \in \mathbb{Z}^d \setminus \{0\}$ such that $C' \cap (C' + v) \neq \emptyset$; i.e., C' and a translate of C' by an integer vector intersect.

Proof. By contradiction; suppose the claim is false. Let R be a large integer number. Consider the family $\mathcal{C}$ of translates of C' by the integer vectors in the cube $[-R, R]^d$: $\mathcal{C} = \{C' + v\colon v \in [-R, R]^d \cap \mathbb{Z}^d\}$, as is indicated in the drawing (C is painted in gray).

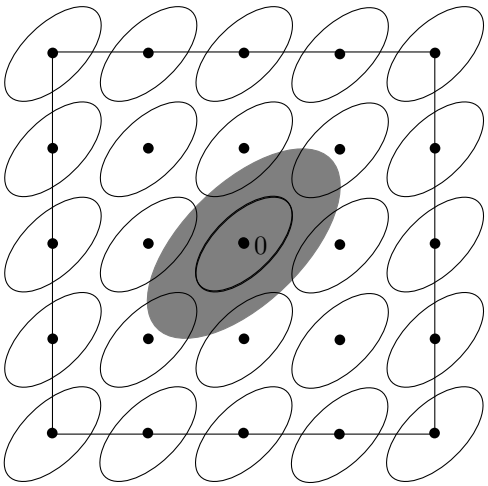

Each such translate is disjoint from C', and thus every two of these translates are disjoint as well. They are all contained in the enlarged cube $K = [-R - D, R + D]^d$, where D denotes the diameter of C'. Hence

$$\mathrm{vol}(K) = (2R + 2D)^d \geq |\mathcal{C}|\,\mathrm{vol}(C') = (2R + 1)^d\,\mathrm{vol}(C'),$$

and

$$\mathrm{vol}(C') \leq \left(1 + \frac{2D - 1}{2R + 1}\right)^d.$$

The expression on the right-hand side is arbitrarily close to 1 for sufficiently large R. On the other hand, $\mathrm{vol}(C') =$

$2^{-d}\,\mathrm{vol}(C) > 1$ is a fixed number exceeding 1 by a certain amount independent of R, a contradiction. The claim thus holds. □

Now let us fix a $v \in \mathbb{Z}^d$ as in the claim and let us choose a point $x \in C' \cap (C' + v)$. Then we have $x - v \in C'$, and since C' is symmetric, we obtain $v - x \in C'$. Since C' is convex, the midpoint of the segment $x(v - x)$ lies in C' too, and so we have $\frac{1}{2}x + \frac{1}{2}(v - x) = \frac{1}{2}v \in C'$. This means that $v \in C$, which proves Minkowski's theorem. □

2.1.2 Example (About a regular forest). Let K be a circle of diameter 26 (meters, say) centered at the origin. Trees of diameter 0.16 grow at each lattice point within K except for the origin, which is where you are standing. Prove that you cannot see outside this miniforest.

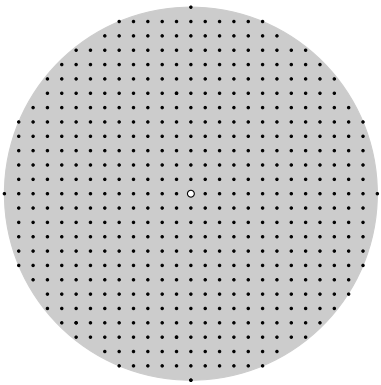

Proof. Suppose than one could see outside along some line ℓ passing through the origin. This means that the strip S of width 0.16 with ℓ as the middle line contains no lattice point in K except for the origin. In other words, the symmetric convex set $C = K \cap S$ contains no lattice points but the origin. But as is easy to calculate, $\mathrm{vol}(C) > 4$, which contradicts Minkowski's theorem. □

2.1.3 Proposition (Approximating an irrational number by a fraction). *Let $\alpha \in (0,1)$ be a real number and N a natural number. Then there exists a pair of natural numbers m, n such that $n \le N$ and*

$$\left|\alpha - \frac{m}{n}\right| < \frac{1}{nN}.$$

This proposition implies that there are infinitely many pairs m, n such that $|\alpha - \frac{m}{n}| < 1/n^2$. This is a basic and well-known result in elementary number theory. It can also be proved using the pigeonhole principle.

Proof of Proposition 2.1.3. Consider the set

$$C = \left\{(x,y) \in \mathbb{R}^2 \colon -N - \tfrac{1}{2} \le x \le N + \tfrac{1}{2},\ |\alpha x - y| < \tfrac{1}{N}\right\}.$$

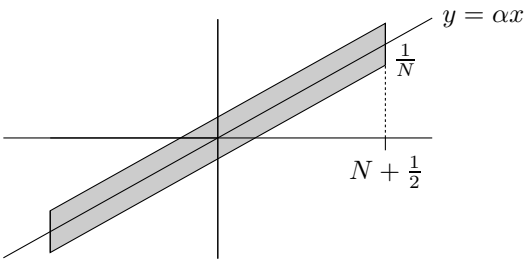

This is a symmetric convex set of area $(2N+1)\frac{2}{N} > 4$, and therefore it contains some nonzero integer lattice point (n, m). By symmetry, we may assume $n > 0$. The definition of C gives $n \le N$ and $|\alpha n - m| < \frac{1}{N}$. In other words, $|\alpha - \frac{m}{n}| < \frac{1}{nN}$. □

2.2 General Lattices

Let $z_1, z_2, \ldots, z_d$ be a d-tuple of linearly independent vectors in $\mathbb{R}^d$. We define the *lattice with basis* $\{z_1, z_2, \ldots, z_d\}$ as the set of all linear combinations of the z_i with integer coefficients; that is,

$$\Lambda = \Lambda(z_1, z_2, \ldots, z_d) = \{i_1 z_1 + i_2 z_2 + \cdots + i_d z_d \colon (i_1, i_2, \ldots, i_d) \in \mathbb{Z}^d\}.$$

Let us remark that this lattice has in general many different bases. For instance, the sets $\{(0,1),(1,0)\}$ and $\{(1,0),(3,1)\}$ are both bases of the "standard" lattice $\mathbb{Z}^2$.

Let us form a $d \times d$ matrix Z with the vectors $z_1, \ldots, z_d$ as columns. We define the *determinant of the lattice* $\Lambda = \Lambda(z_1, z_2, \ldots, z_d)$ as $\det \Lambda = |\det Z|$. Geometrically, $\det \Lambda$ is the volume of the parallelepiped $\{\alpha_1 z_1 + \alpha_2 z_2 + \cdots + \alpha_d z_d \colon \alpha_1, \ldots, \alpha_d \in [0,1]\}$:

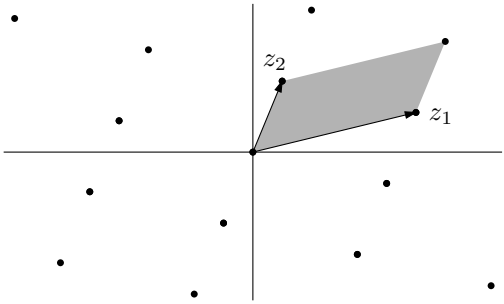

(the proof is left as an exercise). The number $\det \Lambda$ is indeed a property of the lattice Λ (as a point set) and it does not depend on the choice of the basis of Λ (exercise). It is not difficult to show that if Z is the matrix of some basis of Λ, then the matrix of every basis of Λ has the form ZU, where U is an integer matrix with determinant ± 1.

2.2.1 Theorem (Minkowski's theorem for general lattices). *Let Λ be a lattice in $\mathbb{R}^d$, and let $C \subseteq \mathbb{R}^d$ be a symmetric convex set with $\mathrm{vol}(C) > 2^d \det \Lambda$. Then C contains a point of Λ different from 0.*

Proof. Let $\{z_1, \ldots, z_d\}$ be a basis of Λ. We define a linear mapping $f \colon \mathbb{R}^d \to \mathbb{R}^d$ by $f(x_1, x_2, \ldots, x_d) = x_1 z_1 + x_2 z_2 + \cdots + x_d z_d$. Then f is a bijection and $\Lambda = f(\mathbb{Z}^d)$. For any convex set X, we have $\mathrm{vol}(f(X)) = \det(\Lambda)\,\mathrm{vol}(X)$. (Sketch of proof: This holds if X is a cube, and a convex set can be approximated by a disjoint union of sufficiently small cubes with arbitrary precision.) Let us put $C' = f^{-1}(C)$. This is a symmetric convex set with $\mathrm{vol}(C') = \mathrm{vol}(C)/\det \Lambda > 2^d$. Minkowski's theorem

provides a nonzero vector $v \in C' \cap \mathbb{Z}^d$, and $f(v)$ is the desired point as in the theorem. □

A seemingly more general definition of a lattice. What if we consider integer linear combinations of more than d vectors in $\mathbb{R}^d$? Some caution is necessary: If we take $d = 1$ and the vectors $v_1 = (1)$, $v_2 = (\sqrt{2})$, then the integer linear combinations $i_1 v_1 + i_2 v_2$ are dense in the real line (by Proposition 2.1.3), and such a set is not what we would like to call a lattice.

In order to exclude such pathology, we define a *discrete subgroup of* $\mathbb{R}^d$ as a set $\Lambda \subset \mathbb{R}^d$ such that whenever $x, y \in \Lambda$, then also $x - y \in \Lambda$, and such that the distance of any two distinct points of Λ is at least δ, for some fixed positive real number $\delta > 0$.

It can be shown, for instance, that if $v_1, v_2, \ldots, v_n \in \mathbb{R}^d$ are vectors with *rational* coordinates, then the set Λ of all their integer linear combinations is a discrete subgroup of $\mathbb{R}^d$ (exercise). As the following theorem shows, any discrete subgroup of $\mathbb{R}^d$ whose linear span is all of $\mathbb{R}^d$ is a lattice in the sense of the definition given at the beginning of this section.

2.2.2 Theorem (Lattice basis theorem). *Let $\Lambda \subset \mathbb{R}^d$ be a discrete subgroup of $\mathbb{R}^d$ whose linear span is $\mathbb{R}^d$. Then Λ has a basis; that is, there exist d linearly independent vectors $z_1, z_2, \ldots, z_d \in \mathbb{R}^d$ such that $\Lambda = \Lambda(z_1, z_2, \ldots, z_d)$.*

Proof. We proceed by induction. For some i, $1 \le i \le d{+}1$, suppose that linearly independent vectors $z_1, z_2, \ldots, z_{i-1} \in \Lambda$ with the following property have already been constructed. If F_{i-1} denotes the $(i{-}1)$-dimensional subspace spanned by $z_1, \ldots, z_{i-1}$, then all points of Λ lying in F_{i-1} can be written as integer linear combinations of $z_1, \ldots, z_{i-1}$. For $i = d{+}1$, this gives the statement of the theorem.

So consider an $i \le d$. Since Λ generates $\mathbb{R}^d$, there exists a vector $w \in \Lambda$ not lying in the subspace F_{i-1}. Let P be the i-dimensional parallelepiped determined by $z_1, z_2, \ldots, z_{i-1}$ and by w: $P = \{\alpha_1 z_1 + \alpha_2 z_2 + \cdots + \alpha_{i-1} z_{i-1} + \alpha_i w \colon \alpha_1, \ldots, \alpha_i \in [0, 1]\}$. Among all the (finitely

many) points of Λ lying in P but not in F_{i-1}, choose one nearest to F_{i-1} and call it z_i, as in the picture:

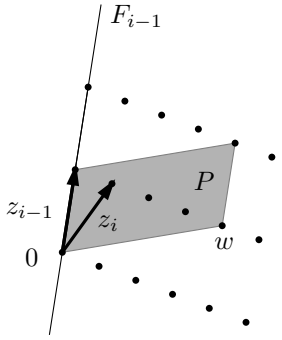

Note that if the points of $\Lambda \cap P$ are written in the form $\alpha_1 z_1 + \alpha_2 z_2 + \cdots + \alpha_{i-1} z_{i-1} + \alpha_i w$, then z_i is one with the smallest positive α_i. It remains to show that $z_1, z_2, \ldots, z_i$ have the required property.

So let $v \in \Lambda$ be a point lying in F_i (the linear span of $z_1, z_2, \ldots, z_i$). We have $v = \sum_{j=1}^{i-1} \beta_j z_j + \beta_i w$ for some real numbers $\beta_1, \ldots, \beta_i$ (since $z_1, z_2, \ldots, z_{i-1}$ and w also span F_i). Further, let us write $z_i = \sum_{j=1}^{i-1} \alpha_j z_j + \alpha_i w$, $0 < \alpha_i \leq 1$. We let $q = \lfloor \beta_i / \alpha_i \rfloor$ and $v' = v - q z_i = \sum_{j=1}^{i-1} \gamma_j z_j + \gamma_i w \in \Lambda$, where $0 \leq \gamma_i = \beta_i - q\alpha_i < \alpha_i$ (the other γ_j could be expressed as well but they don't matter). If $\gamma_i = 0$, then $v' \in \Lambda \cap F_{i-1}$, and we are done by the inductive hypothesis. If $\gamma_i > 0$, then $v'' = v' - \sum_{j=1}^{i-1} \lfloor \gamma_j \rfloor z_j$ is a lattice point in P that is nearer to F_{i-1} than z_i, since the coefficient of w is $\gamma_i < \alpha_i$. This contradiction finishes the induction step. □

Therefore, a lattice can also be defined as a full-dimensional discrete subgroup of $\mathbb{R}^d$.

2.3 An Application in Number Theory

We prove one nontrivial result of elementary number theory. The proof via Minkowski's theorem is one of several possible proofs. Another proof uses the pigeonhole principle in a clever way.

2.3.1 Theorem (Two-square theorem). *Each prime $p \equiv 1 \pmod 4$ can be written as a sum of two squares: $p = a^2 + b^2$, $a, b \in \mathbb{Z}$.*

Let $F = GF(p)$ stand for the field of residue classes modulo p, and let $F^* = F \setminus \{0\}$. An element $a \in F^*$ is called a *quadratic residue* modulo p if there exists an $x \in F^*$ with $x^2 \equiv a \pmod p$. Otherwise, a is a *quadratic nonresidue*.

2.3.2 Lemma. *If p is a prime with $p \equiv 1 \pmod 4$ then -1 is a quadratic residue modulo p.*

Proof. The equation $i^2 = 1$ has two solutions in the field F, namely $i = 1$ and $i = -1$. Hence for any $i \neq \pm 1$ there exists exactly one $j \neq i$ with $ij = 1$ (namely, $j = i^{-1}$, the inverse element in F), and all the elements of $F^* \setminus \{-1, 1\}$ can be divided into pairs such that the product of elements in each pair is 1. Therefore, $(p-1)! = 1 \cdot 2 \cdots (p-1) \equiv -1 \pmod p$.

For a contradiction, suppose that the equation $i^2 = -1$ has no solution in F. Then all the elements of F^* can be divided into pairs such that the product of the elements in each pair is -1. The number of pairs is $(p-1)/2$, which is an even number. Hence $(p-1)! \equiv (-1)^{(p-1)/2} = 1$, a contradiction. □

Proof of Theorem 2.3.1. By the lemma, we can choose a number q such that $q^2 \equiv -1 \pmod p$. Consider the lattice $\Lambda = \Lambda(z_1, z_2)$, where $z_1 = (1, q)$ and $z_2 = (0, p)$. We have $\det \Lambda = p$. We use Minkowski's theorem for general lattices (Theorem 2.2.1) for the disk $C = \{(x, y) \in \mathbb{R}^2 \colon x^2 + y^2 < 2p\}$. The area of C is $2\pi p > 4p = 4 \det \Lambda$, and so C contains a point $(a, b) \in \Lambda \setminus \{0\}$. We have $0 < a^2 + b^2 < 2p$. At the same time, $(a, b) = iz_1 + jz_2$ for some $i, j \in \mathbb{Z}$, which means that $a = i, b = iq + jp$. We calculate $a^2 + b^2 = i^2 + (iq + jp)^2 = i^2 + i^2q^2 + 2iqjp + j^2p^2 \equiv i^2(1 + q^2) \equiv 0 \pmod p$. Therefore $a^2 + b^2 = p$. □

3

Convex Independent Subsets

Here we consider geometric Ramsey-type results about finite point sets in the plane. Ramsey-type theorems are generally statements of the following type: Every sufficiently large structure of a given type contains a "regular" substructure of a prescribed size. In the forthcoming Erdős–Szekeres theorem (Theorem 3.1.3), the "structure of a given type" is simply a finite set of points in general position in $\mathbb{R}^2$, and the "regular substructure" is a set of points forming the vertex set of a convex polygon, as is indicated in the picture:

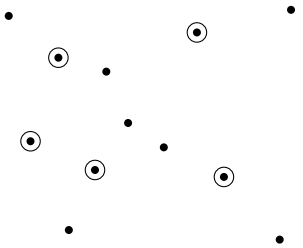

A prototype of Ramsey-type results is Ramsey's theorem itself: *For every choice of natural numbers p, r, n, there exists a natural number N such that whenever X is an N-element set and $c\colon \binom{X}{p} \to \{1, 2, \ldots, r\}$ is an arbitrary coloring of the system of all p-element subsets of X by r colors, then there is an n-element subset $Y \subseteq X$ such that all the p-tuples in $\binom{Y}{p}$ have the same color.* The most famous special case is with

$p = r = 2$, where $\binom{X}{2}$ is interpreted as the edge set of the complete graph K_N on N vertices. Ramsey's theorem asserts that if each of the edges of K_N is colored red or blue, we can always find a complete subgraph on n vertices with all edges red or all edges blue.

Many of the geometric Ramsey-type theorems, including the Erdős–Szekeres theorem, can be derived from Ramsey's theorem. But the quantitative bound for the N in Ramsey's theorem is very large, and consequently, the size of the "regular" configurations guaranteed by proofs via Ramsey's theorem is very small. Other proofs tailored to the particular problems and using more of their geometric structure often yield much better quantitative results.

3.1 The Erdős–Szekeres Theorem

3.1.1 Definition (Convex independent set). *We say that a set $X \subseteq \mathbb{R}^d$ is* convex independent *if for every $x \in X$, we have $x \notin \operatorname{conv}(X \setminus \{x\})$.*

The phrase "in convex position" is sometimes used synonymously with "convex independent." In the plane, a finite convex independent set is the set of vertices of a convex polygon. We will discuss results concerning the occurrence of convex independent subsets in sufficiently large point sets. Here is a simple example of such a statement.

3.1.2 Proposition. *Among any 5 points in the plane in general position (no 3 collinear), we can find 4 points forming a convex independent set.*

Proof. If the convex hull has 4 or 5 vertices, we are done. Otherwise, we have a triangle with two points inside, and the two interior points together with one of the sides of the triangle define a convex quadrilateral. □

Next, we prove a general result.

3.1.3 Theorem (Erdős–Szekeres theorem). *For every natural number k there exists a number $n(k)$ such that any $n(k)$-point set $X \subset \mathbb{R}^2$ in general position contains a k-point convex independent subset.*

First proof (using Ramsey's theorem and Proposition 3.1.2). Color a 4-tuple $T \subset X$ red if its four points are convex independent and blue otherwise. If n is sufficiently large, Ramsey's theorem provides a k-point subset $Y \subset X$ such that all 4-tuples from Y have the same color. But for $k \geq 5$ this color cannot be blue, because any 5 points determine at least one red 4-tuple. Consequently, Y is convex independent, since every 4 of its points are (Carathéodory's theorem). □

Next, we give an inductive proof; it yields an almost tight bound for $n(k)$.

Second proof of the Erdős–Szekeres theorem. In this proof, by a set in general position we mean a set with no 3 points on a common line and no 2 points having the same x-coordinate. The latter can always be achieved by rotating the coordinate system.

Let X be a finite point set in the plane in general position. We call X a *cup* if X is convex independent and its convex hull is bounded from above by a single edge (in other words, if the points of X lie on the graph of a convex function).

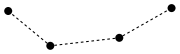

Similarly, we define a *cap*, with a single edge bounding the convex hull from below.

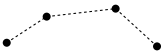

A k-cap is a cap with k points, and similarly for an ℓ-cup.

We define $f(k, \ell)$ as the smallest number N such than any N-point set in general position contains a k-cup or an ℓ-cap. By induction on k and ℓ, we prove the following formula for $f(k, \ell)$:

$$f(k, \ell) \leq \binom{k + \ell - 4}{k - 2} + 1. \tag{3.1}$$

Theorem 3.1.3 clearly follows from this, with $n(k) \leq f(k, k)$. For $k \leq 2$ or $\ell \leq 2$ the formula holds. Thus, let $k, \ell \geq 3$, and consider a set P in

general position with $N = f(k-1,\ell) + f(k,\ell-1)-1$ points. We prove that it contains a k-cup or an ℓ-cap. This will establish the inequality $f(k,\ell) \le f(k-1,\ell) + f(k,\ell-1)-1$, and then (3.1) follows by induction; we leave the simple manipulation of binomial coefficients to the reader.

Suppose that there is no ℓ-cap in X. Let $E \subseteq X$ be the set of points $p \in X$ such that X contains a $(k-1)$-cup ending with p.

We have $|E| \ge N - f(k-1,\ell)+1 = f(k,\ell-1)$, because $X \setminus E$ contains no $(k-1)$-cup and so $|X \setminus E| < f(k-1,\ell)$.

Either the set E contains a k-cup, and then we are done, or there is an $(\ell-1)$-cap. The first point p of such an $(\ell-1)$-cap is, by the definition of E, the last point of some $(k-1)$-cup in X, and in this situation, either the cup or the cap can be extended by one point:

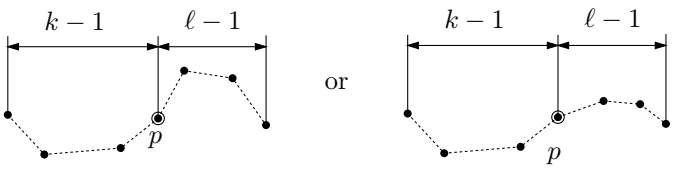

This finishes the inductive step. □

A lower bound for sets without k-cups and ℓ-caps. Interestingly, the bound for $f(k,\ell)$ proved above is tight, not only asymptotically but exactly! This means, in particular, that there are n-point planar sets in general position where any convex independent subset has at most $O(\log n)$ points, which is somewhat surprising at first sight.

An example of a set $X_{k,\ell}$ of $\binom{k+\ell-4}{k-2}$ points in general position with no k-cup and no ℓ-cap can be constructed, again by induction on $k+\ell$. If $k \le 2$ or $\ell \le 2$, then $X_{k,\ell}$ can be taken as a one-point set.

Supposing both $k \ge 3$ and $\ell \ge 3$, the set $X_{k,\ell}$ is obtained from the sets $L = X_{k-1,\ell}$ and $R = X_{k,\ell-1}$ according to the following picture:

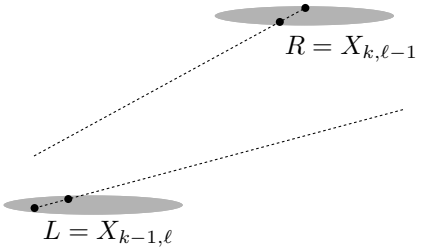

The set L is placed to the left of R in such a way that all lines determined by pairs of points in L go below R and all lines determined by pairs of points of R go above L.

Consider a cup C in the set $X_{k,\ell}$ thus constructed. If $C \cap L = \emptyset$, then $|C| \le k-1$ by the assumption on R. If $C \cap L \ne \emptyset$, then C has at most 1 point in R, and since no cup in L has more than $k-2$ points, we get $|C| \le k-1$ as well. The argument for caps is symmetric.

We have $|X_{k,\ell}| = |X_{k-1,\ell}| + |X_{k,\ell-1}|$, and the formula for $|X_{k,\ell}|$ follows by induction; the calculation is almost the same as in the previous proof. □

Determining the exact value of $n(k)$ in the Erdős–Szekeres theorem is much more challenging. Here are the best known bounds:

$$2^{k-2} + 1 \le n(k) \le \binom{2k-5}{k-2} + 2.$$

The upper bound is a small improvement over the bound $f(k,k)$ derived above. The lower bound results from an inductive construction slightly more complicated than that of $X_{k,\ell}$.

3.2 Horton Sets

Let X be a set in $\mathbb{R}^d$. A k-point set $Y \subseteq X$ is called a *k-hole in X* if Y is convex independent and $\mathrm{conv}(Y) \cap X = Y$. In the plane, Y determines a convex k-gon with no points of X inside. Erdős raised the question about the rather natural strengthening of the Erdős–Szekeres theorem:

Is it true that for every k there exists an $n(k)$ such that any $n(k)$-point set in the plane in general position has a k-hole?

A construction due to Horton, whose streamlined version we present below, shows that this is *false* for $k \geq 7$: There are arbitrarily large sets without a 7-hole. On the other hand, a positive result holds for $k \leq 6$. The result for $k \leq 5$ is not hard and we prove it below. On the other hand, for $k = 6$ the problem had been famous and open for many years, and only recently it has been solved (independently by two authors).

3.2.1 Proposition (The existence of a 5-hole). *Every sufficiently large planar point set in general position contains a 5-hole.*

Proof. By the Erdős–Szekeres theorem, we may assume that there exists a 6-point convex independent subset of our set X. Consider a 6-point convex independent subset $H \subseteq X$ with the smallest possible $|X \cap \mathrm{conv}(H)|$. Let $I = \mathrm{conv}(H) \cap (X \setminus H)$ be the points inside the convex hull of H.

- If $I = \emptyset$, we have a 6-hole.
- If there is one point x in I, we consider a diagonal that partitions the hexagon into two quadrilaterals:

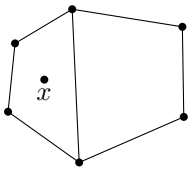

  The point x lies in one of these quadrilaterals, and the vertices of the other quadrilateral together with x form a 5-hole.
- If $|I| \geq 2$, we choose an edge xy of $\mathrm{conv}(I)$. Let γ be an open half-plane bounded by the line xy and containing no points of I (it is determined uniquely unless $|I| = 2$).

If $|\gamma \cap H| \geq 3$, we get a 5-hole formed by x, y, and 3 points of $\gamma \cap H$. For $|\gamma \cap H| \leq 2$, we have one of the two cases indicated in the following picture:

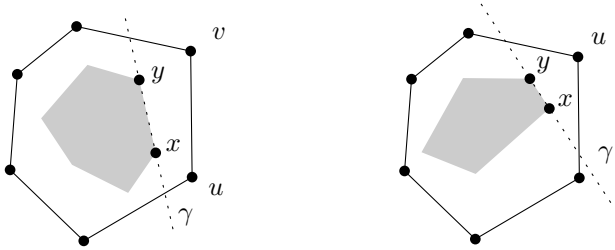

By replacing u and v by x and y in the left situation, or u by x in the right situation, we obtain a 6-point convex independent set having fewer points inside than H, which is a contradiction. □

3.2.2 Theorem (Seven-hole theorem). *There exist arbitrarily large finite sets in the plane in general position without a 7-hole.*

The sets constructed in the proof have other interesting properties as well.

Definitions. Let X and Y be finite sets in the plane. We say that X is *high above* Y (and that Y is *deep below* X) if the following hold:

(i) No line determined by two points of $X \cup Y$ is vertical.

(ii) Each line determined by two points of X lies above all the points of Y.

(iii) Each line determined by two points of Y lies below all the points of X.

For a set $X = \{x_1, x_2, \ldots, x_n\}$, with no two points having equal x-coordinates and with notation chosen so that the x-coordinates of the x_i increase with i, we define the sets $X_0 = \{x_2, x_4, \ldots\}$ (consisting of the

points with even indices) and $X_1 = \{x_1, x_3, \ldots\}$ (consisting of the points with odd indices).

A finite set $H \subset \mathbb{R}^2$ is a *Horton set* if $|H| \leq 1$, or the following conditions hold: $|H| > 1$, both H_0 and H_1 are Horton sets, and H_1 lies high above H_0 or H_0 lies high above H_1.

3.2.3 Lemma. *For every $n \geq 1$, an n-point Horton set exists.*

Proof. We note that one can produce a smaller Horton set from a larger one by deleting points from the right. We construct $H^{(k)}$, a Horton set of size 2^k, by induction.

We define $H^{(0)}$ as the point $(0, 0)$. Suppose that we can construct a Horton set $H^{(k)}$ with 2^k points whose x-coordinates are $0, 1, \ldots, 2^k-1$. The induction step goes as follows.

Let $A = 2H^{(k)}$ (i.e., $H^{(k)}$ expanded twice), and $B = A+(1, h_k)$, where h_k is a sufficiently large number. We set $H^{(k+1)} = A \cup B$. It is easily seen that if h_k is large enough, B lies high above A, and so $H^{(k+1)}$ is Horton as well. The set $H^{(3)}$ looks like this:

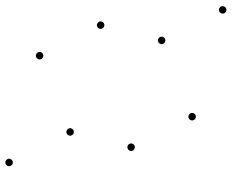

□

Closedness from above and from below. A set X in $\mathbb{R}^2$ is *r-closed from above* if for any r-cup in X there exists a point in X lying above the r-cup (i.e., above the bottom part of its convex hull).

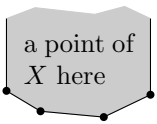

$r = 4$

Similarly, we define a set *r-closed from below* using r-caps.

3.2.4 Lemma. *Every Horton set is both* 4*-closed from above and* 4*-closed from below.*

Proof. We proceed by induction on the size of the Horton set. Let H be a Horton set, and assume that H_0 lies deep below H_1 (the other possible case is analogous). Let $C \subseteq H$ be a 4-cup.

If $C \subseteq H_0$ or $C \subseteq H_1$, then a point closing C from above exists by the inductive hypothesis. Thus, let $C \cap H_0 \neq \emptyset \neq C \cap H_1$.

The cup C may have at most 2 points in H_1 (the upper part): If there were 3 points, say a, b, c (in left-to-right order), then H_0 lies below the lines ab and bc, and so the remaining point of C, which was supposed to lie in H_0, cannot form a cup with $\{a, b, c\}$:

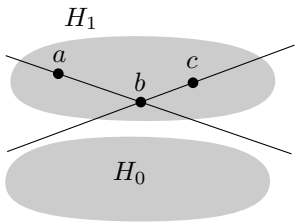

This means that C has at least 2 points, a and b, in the lower part H_0. Since the points of H_0 and H_1 alternate along the x-axis, there is a point $c \in H_1$ between a and b in the ordering by x-coordinates. This c is above the segment ab, and so it closes the cup C from above. We argue similarly for a 4-cap. □

3.2.5 Proposition. *No Horton set contains a* 7*-hole.*

Proof. (Very similar to the previous one.) For contradiction, suppose there is a 7-hole X in the considered Horton set H. If $X \subseteq H_0$ or $X \subseteq H_1$, we use induction. Otherwise, we select the part (H_0 or H_1) containing the larger portion of X; this has at least 4 points of X. If this part is, say, H_0, and it lies deep below H_1, these 4 points must form a cup in H_0, for if some 3 of them were a cap, no point of H_1 could complete them

to a convex independent set. By Lemma 3.2.4, H_0 (being a Horton set) contains a point closing the 4-cup from above. Such a point must be contained in the convex hull of the 7-hole X, a contradiction. □

4

Incidence Problems

In this chapter we study a very natural problem of combinatorial geometry: the maximum possible number of incidences between m points and n lines in the plane. In addition to its mathematical appeal, this problem and its relatives are significant in the analysis of several basic geometric algorithms. In the proofs we encounter number-theoretic arguments, results about graph drawing, and the probabilistic method.

4.1 Formulation

Point–line incidences. Consider a set P of m points and a set L of n lines in the plane. What is the maximum possible number of their *incidences*, i.e., pairs (p, ℓ) such that $p \in P$, $\ell \in L$, and p lies on ℓ? We denote the number of incidences for specific P and L by $I(P, L)$, and we let $I(m, n)$ be the maximum of $I(P, L)$ over all choices of an m-element P and an n-element L. For example, the following picture illustrates that $I(3, 3) \geq 6$,

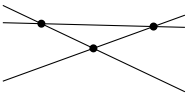

and it is not hard to see that actually $I(3, 3) = 6$.

A trivial upper bound is $I(m,n) \le mn$, but it it can never be attained unless $m = 1$ or $n = 1$. In fact, if m has a similar order of magnitude as n then $I(m,n)$ is asymptotically much smaller than mn. The order of magnitude is known exactly:

4.1.1 Theorem (Szemerédi–Trotter theorem). *For all $m, n \ge 1$, we have $I(m,n) = O(m^{2/3}n^{2/3}+m+n)$, and this bound is asymptotically tight.*

We will mostly consider only the most interesting case $m = n$. The general case needs no new ideas but only a little more complicated calculation.

Of course, the problem of point–line incidences can be generalized in many ways. We can consider incidences between points and hyperplanes in higher dimensions, or between points in the plane and some family of curves, and so on. A particularly interesting case is that of points and unit circles, which is closely related to counting unit distances.

Unit distances and distinct distances. Let $U(n)$ denote the maximum possible number of pairs of points with unit distance in an n-point set in the plane. For $n \le 3$ we have $U(n) = \binom{n}{2}$ (all distances can be 1), but already for $n = 4$ at most 5 of the 6 distances can be 1; i.e., $U(4) = 5$:

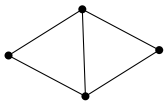

We are interested in the asymptotic behavior of the function $U(n)$ for $n \to \infty$.

This can also be reformulated as an incidence problem. Namely, consider an n-point set P and draw a unit circle around each point of p, thereby obtaining a set C of n unit circles. Each pair of points at unit distance contributes two point–circle incidences, and hence $U(n) \le \frac{1}{2}I_{1\mathrm{circ}}(n,n)$, where $I_{1\mathrm{circ}}(m,n)$ denotes the maximum possible number of incidences between m points and n unit circles.

Unlike the case of point–line incidences, the correct order of magnitude of $U(n)$ is not known. An upper bound of $O(n^{4/3})$ can be obtained by modifying proofs of the Szemerédi–Trotter theorem. But the best known lower bound is $U(n) \geq n^{1+c_1/\log\log n}$, for some positive constant c_1; this is superlinear in n but grows more slowly than $n^{1+\varepsilon}$ for every fixed $\varepsilon > 0$.

A related quantity is the minimum possible number of distinct distances determined by n points in the plane; formally,

$$g(n) = \min_{P\subset\mathbb{R}^2:\,|P|=n} |\{\mathrm{dist}(x,y)\colon x,y\in P\}|.$$

Clearly, $g(n) \geq \binom{n}{2}/U(n)$, and so the bound $U(n) = O(n^{4/3})$ mentioned above gives $g(n) = \Omega(n^{2/3})$. This has been improved several times, and the current best lower bound is approximately $\Omega(n^{0.863})$. The best known upper bound is $O(n/\sqrt{\log n})$.

4.2 Lower Bounds: Incidences and Unit Distances

4.2.1 Proposition (Many point–line incidences). *We have $I(n,n) = \Omega(n^{4/3})$, and so the upper bound for the maximum number of incidences of n points and n lines in the plane in the Szemerédi–Trotter theorem is asymptotically optimal.*

It is not easy to come up with good constructions "by hand." Small cases do not seem to be helpful for discovering a general pattern. Surprisingly, an asymptotically optimal construction is quite simple. The appropriate lower bound for $I(m,n)$ with $n \neq m$ can be obtained similarly.

Proof. For simplicity, we suppose that $n = 4k^3$ for a natural number k. For the point set P, we choose the $k \times 4k^2$ grid; i.e., we set $P = \{(i,j)\colon i = 0,1,2,\ldots,k{-}1,\ j = 0,1,\ldots,4k^2{-}1\}$. The set L consists of all the lines with equations $y = ax + b$, where $a = 0,1,\ldots,2k{-}1$ and $b = 0,1,\ldots,2k^2{-}1$. These are n lines, as it should be. For $x \in [0,k)$,

we have $ax + b < ak + b < 2k^2 + 2k^2 = 4k^2$. Therefore, for each $i = 0, 1, \ldots, k-1$, each line of L contains a point of P with the x-coordinate equal to i, and so $I(P, L) \geq k \cdot |L| = \Omega(n^{4/3})$. □

Next, we consider unit distances, where the construction is equally simple but the analysis uses considerable number-theoretic tools.

4.2.2 Theorem (Many unit distances). *For all $n \geq 2$, there exist configurations of n points in the plane determining at least $n^{1+c_1/\log\log n}$ unit distances, with a positive constant c_1.*

A configuration with the asymptotically largest known number of unit distances is a $\sqrt{n} \times \sqrt{n}$ regular grid with a suitably chosen step. Here unit distances are related to the number of possible representations of an integer as a sum of two squares. We begin with the following claim:

4.2.3 Lemma. *Let $p_1 < p_2 < \cdots < p_r$ be primes of the form $4k+1$, and put $M = p_1 p_2 \cdots p_r$. Then M can be expressed as a sum of two squares of integers in at least 2^r ways.*

Proof. As we know from Theorem 2.3.1, each p_j can be written as a sum of two squares: $p_j = a_j^2 + b_j^2$. In the sequel, we work with the ring $\mathbb{Z}[\mathrm{i}]$, the so-called Gaussian integers, consisting of all complex numbers $u + \mathrm{i}v$, where $u, v \in \mathbb{Z}$. We use the fact that each element of $\mathbb{Z}[\mathrm{i}]$ can be uniquely factored into primes. From algebra, we recall that a *prime* in the ring $\mathbb{Z}[\mathrm{i}]$ is an element $\gamma \in \mathbb{Z}[\mathrm{i}]$ such that whenever $\gamma = \gamma_1\gamma_2$ with $\gamma_1, \gamma_2 \in \mathbb{Z}[\mathrm{i}]$, then $|\gamma_1| = 1$ or $|\gamma_2| = 1$. Both existence and uniqueness of prime factorization follows from the fact that $\mathbb{Z}[\mathrm{i}]$ is a Euclidean ring (see an introductory course on algebra for an explanation of these notions).

Let us put $\alpha_j = a_j + \mathrm{i}\,b_j$, and let $\bar\alpha_j = a_j - \mathrm{i}\,b_j$ bethe complex conjugate of α_j. We have $\alpha_j\bar\alpha_j = (a_j + \mathrm{i}\,b_j)(a_j - \mathrm{i}\,b_j) = a_j^2 + b_j^2 = p_j$. Let us choose an arbitrary subset $J \subseteq I = \{1, 2, \ldots, r\}$ and define $A_J + \mathrm{i}B_J = \left(\prod_{j\in J}\alpha_j\right)\left(\prod_{j\in I\setminus J}\bar\alpha_j\right)$. Then $A_J - \mathrm{i}B_J = \left(\prod_{j\in J}\bar\alpha_j\right)\left(\prod_{j\in I\setminus J}\alpha_j\right)$, and hence $M = (A_J + \mathrm{i}B_J)(A_J - \mathrm{i}B_J) = A_J^2 + B_J^2$. This gives one expression

of the number M as a sum of two squares. It remains to prove that for two sets $J \neq J'$, $A_J + \mathrm{i}B_J \neq A_{J'} + \mathrm{i}B_{J'}$. To this end, it suffices to show that all the α_j and $\bar{\alpha}_j$ are primes in $\mathbb{Z}[\mathrm{i}]$. Then the numbers $A_J + \mathrm{i}B_J$ and $A_{J'} + \mathrm{i}B_{J'}$ are distinct, since they have distinct prime factorizations. (No α_j or $\bar{\alpha}_j$ can be obtained from another one by multiplying it by a unit of the ring $\mathbb{Z}[\mathrm{i}]$: The units are only the elements $1, -1, \mathrm{i}$, and $-\mathrm{i}$.)

So suppose that $\alpha_j = \gamma_1\gamma_2$, $\gamma_1, \gamma_2 \in \mathbb{Z}[\mathrm{i}]$. We have $p_j = \alpha_j\bar{\alpha}_j = \gamma_1\gamma_2\bar{\gamma}_1\bar{\gamma}_2 = |\gamma_1|^2|\gamma_2|^2$. Now, $|\gamma_1|^2$ and $|\gamma_2|^2$ are both integers, and since p_j is a prime, we get that $|\gamma_1| = 1$ or $|\gamma_2| = 1$. □

Next, we need to know that the primes of the form $4k+1$ are sufficiently dense. First we recall the well-known *prime number theorem*: If $\pi(n)$ denotes the number of primes not exceeding n, then

$$\pi(n) = (1 + o(1))\frac{n}{\ln n} \quad \text{as } n \to \infty.$$

Proofs of this fact are quite complicated; on the other hand, it is not so hard to prove weaker bounds $cn/\log n < \pi(n) < Cn/\log n$ for suitable positive constants c, C.

We consider primes in the arithmetic progression $1, 5, 9, \ldots, 4k+1, \ldots$. A famous theorem of Dirichlet asserts that every arithmetic progression contains infinitely many primes unless this is impossible for a trivial reason, namely, unless all the terms have a nontrivial common divisor. The following theorem is still stronger:

4.2.4 Theorem. *Let d and a be relatively prime natural numbers, and let $\pi_{d,a}(n)$ be the number of primes of the form $a + kd$ ($k = 0, 1, 2, \ldots$) not exceeding n. We have*

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

where φ denotes the Euler function: $\varphi(d)$is the number of integers between 1 and d that are relatively prime to d.

For every $d \geq 2$, there are $\varphi(d)$ residue classes modulo d that can possibly contain primes. The theorem shows that the primes are quite uniformly distributed among these residue classes.

The proof of the theorem is not simple, and we omit it, but it is very nice, and we can only recommend to the reader to look it up in a textbook on number theory.

Proof of the lower bound for unit distances (Theorem 4.2.2). Let us suppose that n is a square. For the set P we choose the points of the $\sqrt{n} \times \sqrt{n}$ grid with step $1/\sqrt{M}$, where M is the product of the first $r-1$ primes of the form $4k+1$, and r is chosen as the largest number such that $M \le \frac{n}{4}$.

It is easy to see that each point of the grid participates in at least as many unit distances as there are representations of M as a sum of two squares of nonnegative integers. Since one representation by a sum of two squares of nonnegative integers corresponds to at most 4 representations by a sum of two squares of arbitrary integers (the signs can be chosen in 4 ways), we have at least $2^{r-1}/16$ unit distances by Lemma 4.2.3.

By the choice of r, we have $4p_1p_2\cdots p_{r-1} \le n < 4p_1p_2\cdots p_r$, and hence $2^r \le n$ and $p_r > (\frac{n}{4})^{1/r}$. Further, we obtain, by Theorem 4.2.4, $r = \pi_{4,1}(p_r) \ge (\frac{1}{2} - o(1))p_r/\log p_r \ge \sqrt{p_r} \ge n^{1/3r}$ for sufficientlylarge n, and thus $r^{3r} \ge n$. Taking logarithms, we have $3r\log r \ge \log n$, and hence $r \ge \log n/(3\log r) \ge \log n/(3\log\log n)$. The number of unit distances is at least $n\,2^{r-4} \ge n^{1+c_1/\log\log n}$, as Theorem 4.2.2 claims. Let us remark that for sufficiently large n the constant c_1 can be made as close to 1 as desired. □

4.3 Point–Line Incidences via Crossing Numbers

Here we present a very simple proof of the Szemerédi–Trotter theorem based on a result concerning graph drawing. We need the notion of the crossing number of a graph G; this is the minimum possible number of edge crossings in a drawing of G. To make this rigorous, let us first recall a formal definition of a drawing.

An *arc* is the image of a continuous injective map $[0,1] \to \mathbb{R}^2$. A *drawing* of a graph G is a mapping that assigns to each vertex of G a point in the plane (distinct vertices being assigned distinct points) and

to each edge of G an arc connecting the corresponding two (images of) vertices and not incident to any other vertex. We do not insist that the drawing be planar, so the arcs are allowed to cross. A *crossing* is a point common to at least two arcs but distinct from all vertices.

In this section we will actually deal only with drawings where each edge is represented by a straight segment.

Let G be a graph (or multigraph). The *crossing number* of a drawing of G in the plane is the number of crossings in the considered drawing, where a crossing incident to $k \geq 2$ edges is counted $\binom{k}{2}$ times. So a drawing is planar if and only if its crossing number is 0. The *crossing number* of the graph G is the smallest possible crossing number of a drawing of G; we denote it by $\mathrm{cr}(G)$. For example, $\mathrm{cr}(K_5) = 1$:

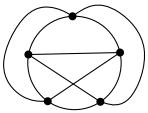

As is well known, for $n > 2$, a planar graph with n vertices has at most $3n-6$ edges. This can be rephrased as follows: If the number of edges is at least $3n-5$ then $\mathrm{cr}(G) > 0$. The following theorem can be viewed as a generalization of this fact.

4.3.1 Theorem (Crossing number theorem). *Let $G = (V, E)$ be a simple graph (no multiple edges). Then*

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

(the constant $\frac{1}{64}$ can be improved by a more careful calculation).

The lower bound in this theorem is asymptotically tight; i.e., there exist graphs with n vertices, m edges, and crossing number $O(m^3/n^2)$ (exercise). The assumption that the graph is simple cannot be omitted.

For a proof of this theorem, we need a simple lemma:

4.3.2 Lemma. *The crossing number of any simple graph $G = (V, E)$ is at least $|E| - 3|V|$.*

Proof. If $|E| > 3|V|$ and some drawing of the graph had fewer than $|E| - 3|V|$ crossings, then we could delete one edge from each crossing and obtain a planar graph with more than $3|V|$ edges. □

Proof of Theorem 4.3.1. Consider some drawing of a graph $G = (V, E)$ with n vertices, m edges, and crossing number x. We may and will assume that no two edges sharing a common vertex cross in the drawing. If we have a straight-edge drawing, as will be the case in the proof of the Szemerédi–Trotter theorem, then this is satisfied automatically, and for curvilinear drawings we can eliminate crossings of neighboring edges by modifying the drawing:

We may assume $m \geq 4n$, for otherwise, the claimed bound is negative. Let $p \in (0, 1)$ be a parameter; later on we set it to a suitable value. We choose a random subset $V' \subseteq V$ by including each vertex $v \in V$ into V' independently with probability p. Let G' be the subgraph of G induced by the subset V'. Put $n' = |V'|$, $m' = |E(G')|$, and let x' be the crossing number of the graph G' in the drawing "inherited" from the considered drawing of G. The expectation of n' is $\mathbf{E}[n'] = np$. The probability that a given edge appears in $E(G')$ is p^2, and hence $\mathbf{E}[m'] = mp^2$, and similarly we get $\mathbf{E}[x'] = xp^4$. At the same time, by Lemma 4.3.2 we always have $x' \geq m' - 3n'$, and so this relation holds for the expectations as well: $\mathbf{E}[x'] \geq \mathbf{E}[m'] - 3\mathbf{E}[n']$.

So we have $xp^4 \geq mp^2 - 3np$. Setting $p = \frac{4n}{m}$ (which is at most 1, since we assume $m \geq 4n$), we calculate that

$$x \geq \frac{1}{64}\,\frac{m^3}{n^2}\,.$$

The crossing number theorem is proved. □

Proof of the Szemerédi–Trotter theorem (Theorem 4.1.1). We consider a set P of m points and a set L of n lines in the plane realizing the maximum number of incidences $I(m,n)$. We define a certain topological graph $G = (V,E)$, that is, a graph together with its drawing in the plane. Each point $p \in P$ becomes a vertex of G, and two points $p, q \in P$ are connected by an edge if they lie on a common line $\ell \in L$ next to one another. So we have a drawing of G where the edges are straight segments. This is illustrated below, with G drawn thick:

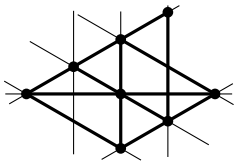

If a line $\ell \in L$ contains $k \geq 1$ points of P, then it contributes $k-1$ edges to E, and hence $I(m,n) = |E| + n$. Since the edges are parts of the n lines, at most $\binom{n}{2}$ pairs may cross: $\mathrm{cr}(G) \leq \binom{n}{2}$. On the other hand, from the crossing number theorem we get $\mathrm{cr}(G) \geq \frac{1}{64} \cdot |E|^3/m^2 - m$. So $\frac{1}{64} \cdot |E|^3/m^2 - m \leq \mathrm{cr}(G) \leq \binom{n}{2}$, and a calculation gives $|E| = O(n^{2/3}m^{2/3} + m)$. This proves the Szemerédi–Trotter theorem. □

The best known upper bound on the number of unit distances, $U(n) = O(n^{4/3})$, can be proved along similar lines (try it!).

5

Convex Polytopes

Convex polytopes are convex hulls of finite point sets in $\mathbb{R}^d$. They constitute the most important class of convex sets with an enormous number of applications and connections.

Three-dimensional convex polytopes, especially the regular ones, have been fascinating people since the antiquity. Their investigation was one of the main sources of the theory of planar graphs, and thanks to this well-developed theory they are quite well understood. But convex polytopes in dimension 4 and higher are considerably more challenging, and a surprisingly deep theory, mainly of algebraic nature, was developed in attempts to understand their structure.

A strong motivation for the study of convex polytopes comes from practically significant areas such as combinatorial optimization, linear programming, and computational geometry. Let us look at a simple example illustrating how polytopes can be associated with combinatorial objects. The 3-dimensional polytope in the picture

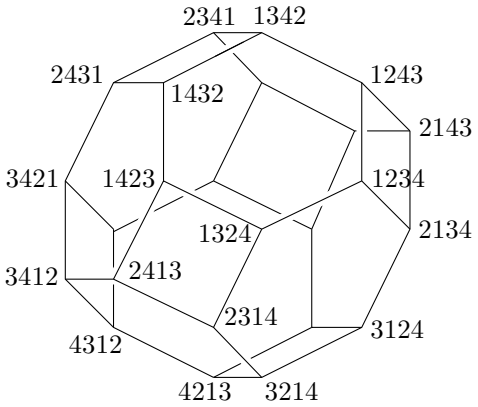

is called the *permutahedron*. Although it is 3-dimensional, it is most naturally defined as a subset of $\mathbb{R}^4$, namely, the convex hull of the 24 vectors obtained by permuting the coordinates of the vector $(1, 2, 3, 4)$ in all possible ways. In the picture, the (visible) vertices are labeled by the corresponding permutations. Similarly, the d-dimensional permutahedron is the convex hull of the $(d+1)!$ vectors in $\mathbb{R}^{d+1}$ arising by permuting the coordinates of $(1, 2, \ldots, d+1)$. One can observe that the edges of the polytope connect exactly pairs of permutations differing by a transposition of two adjacent numbers, and a closer examination reveals other connections between the structure of the permutahedron and properties of permutations.

There are many other, more sophisticated, examples of convex polytopes assigned to combinatorial and geometric objects such as graphs, partially ordered sets, classes of metric spaces, or triangulations of a given point set. In many cases, such convex polytopes are a key tool for proving hard theorems about the original objects or for obtaining efficient algorithms.

5.1 Geometric Duality

First we discuss geometric duality, a simple technical tool indispensable in the study of convex polytopes and handy in many other situations. We begin with a simple motivating question.

How can we visualize the set of all lines intersecting a convex pentagon as in the picture?

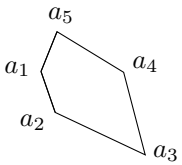

A suitable way is provided by line–point duality.

5.1.1 Definition (Duality transform). *The (geometric)* duality transform *is a mapping denoted by $\mathcal{D}_0$. To a point $a \in \mathbb{R}^d \setminus \{0\}$ it assigns the hyperplane*

$$\mathcal{D}_0(a) = \{x \in \mathbb{R}^d \colon \langle a, x\rangle = 1\},$$

and to a hyperplane h not passing through the origin, which can be uniquely written in the form $h = \{x \in \mathbb{R}^d \colon \langle a, x\rangle = 1\}$, it assigns the point $\mathcal{D}_0(h) = a \in \mathbb{R}^d \setminus \{0\}$.

Here is the geometric meaning of the duality transform. If a is a point at distance δ from 0, then $\mathcal{D}_0(a)$ is the hyperplane perpendicular to the line $0a$ and intersecting that line at distance $\frac{1}{\delta}$ from 0, in the direction from 0 towards a.

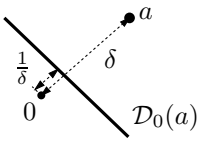

A nice interpretation of duality is obtained by working in $\mathbb{R}^{d+1}$ and identifying the "primal" $\mathbb{R}^d$ with the hyperplane $\pi = \{x \in \mathbb{R}^{d+1} \colon x_{d+1} = 1\}$ and the "dual" $\mathbb{R}^d$ with the hyperplane $\rho = \{x \in \mathbb{R}^{d+1} \colon x_{d+1} = -1\}$. The hyperplane dual to a point $a \in \pi$ is produced as follows: We construct the hyperplane in $\mathbb{R}^{d+1}$ perpendicular to $0a$ and containing 0, and we intersect it with ρ. Here is an illustration for $d = 2$:

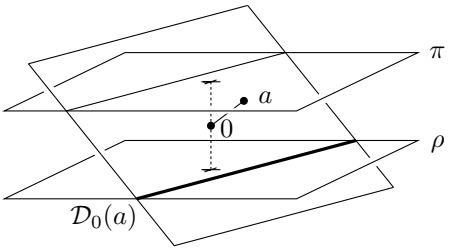

In this way, the duality $\mathcal{D}_0$ can be naturally extended to k-flats in $\mathbb{R}^d$, whose duals are $(d-k-1)$-flats. Namely, given a k-flat $f \subset \pi$, we consider the $(k+1)$-flat F through 0 and f, we construct the orthogonal complement of F, and we intersect it with ρ, obtaining $\mathcal{D}_0(f)$.

Let us consider the pentagon drawn above and place it so that the origin lies in the interior. Let $v_i = \mathcal{D}_0(\ell_i)$, where ℓ_i is the line containing the side $a_i a_{i+1}$. Then the points dual to the lines intersecting the pentagon $a_1 a_2 \ldots a_5$ fill exactly the exterior of the convex pentagon $v_1 v_2 \ldots v_5$:

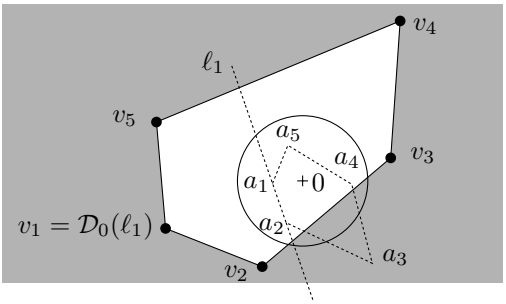

This follows easily from the properties of duality listed below (of course, there is nothing special about a pentagon here). Thus, the considered set of lines can be nicely described in the dual plane. A similar passage from lines to points or back is useful in many geometric or computational problems.

Properties of the duality transform. Let p be a point of $\mathbb{R}^d$ distinct from the origin and let h be a hyperplane in $\mathbb{R}^d$ not containing the origin. Let h^- stand for the closed half-space bounded by h and containing the origin, while h^+ denotes the other closed half-space bounded by h. That is, if $h = \{x \in \mathbb{R}^d\colon \langle a, x\rangle = 1\}$, then $h^- = \{x \in \mathbb{R}^d\colon \langle a, x\rangle \le 1\}$.

5.1.2 Lemma (Duality preserves incidences).

(i) *$p \in h$ if and only if $\mathcal{D}_0(h) \in \mathcal{D}_0(p)$.*

(ii) *$p \in h^-$ if and only if $\mathcal{D}_0(h) \in \mathcal{D}_0(p)^-$.*

Proof. (i) Let $h = \{x \in \mathbb{R}^d \colon \langle a, x\rangle = 1\}$. Then $p \in h$ means $\langle a, p\rangle = 1$. Now, $\mathcal{D}_0(h)$ is the point a, and $\mathcal{D}_0(p)$ is the hyperplane $\{y \in \mathbb{R}^d \colon \langle y, p\rangle = 1\}$, and hence $\mathcal{D}_0(h) = a \in \mathcal{D}_0(p)$ also means just $\langle a, p\rangle = 1$. Part (ii) is proved similarly. □

5.1.3 Definition (Dual set). *For a set $X \subseteq \mathbb{R}^d$, we define the set* dual to *X, denoted by X^*, as follows:*

$$X^* = \{y \in \mathbb{R}^d \colon \langle x, y\rangle \le 1 \text{ for all } x \in X\}\,.$$

Another common name used for the duality is *polarity*; the dual set would then be called the *polar set*. Sometimes it is denoted by X°.

Geometrically, X^* is the intersection of all half-spaces of the form $\mathcal{D}_0(x)^-$ with $x \in X$. Or in other words, X^* consists of the origin plus all points y such that $X \subseteq \mathcal{D}_0(y)^-$. For example, if X is the pentagon $a_1a_2\dots a_5$ drawn above, then X^* is the pentagon $v_1v_2\dots v_5$.

For any set X, the set X^* is obviously closed and convex and contains the origin. Using the separation theorem (Theorem 1.2.4), it is easily shown that for any set $X \subseteq \mathbb{R}^d$, the set $(X^*)^*$ is the closure $\operatorname{conv}(X \cup \{0\})$. In particular, for a closed convex set containing the origin we have $(X^*)^* = X$ (exercise).

For a hyperplane h, the dual set h^* is different from the point $\mathcal{D}_0(h)$.

For readers familiar with the duality of planar graphs, let us remark that it is closely related to the geometric duality applied to convex polytopes in $\mathbb{R}^3$. For example, the next drawing illustrates a planar graph and its dual graph (dashed):

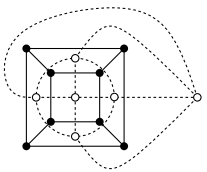

Later we will see that these are graphs of the 3-dimensional cube and of the regular octahedron, which are polytopes dual to each other in the sense defined above. A similar relation holds for all 3-dimensional polytopes and their graphs.

Other variants of duality. The duality transform $\mathcal{D}_0$ defined above is just one of a class of geometric transforms with similar properties. For some purposes, other such transforms (dualities) are more convenient. A particularly important duality, denoted by $\mathcal{D}$, corresponds to moving the origin to the "minus infinity" of the x_d-axis (the x_d-axis is considered vertical). A formal definition is as follows.

5.1.4 Definition (Another duality). *A nonvertical hyperplane h can be uniquely written in the form $h = \{x \in \mathbb{R}^d\colon x_d = a_1x_1 + \cdots + a_{d-1}x_{d-1} - a_d\}$. We set $\mathcal{D}(h) = (a_1, \ldots, a_{d-1}, a_d)$. Conversely, the point a maps back to h.*

The property (i) of Lemma 5.1.2 holds for this $\mathcal{D}$, and an analogue of (ii) is:

(ii′) A point p lies above a hyperplane h if and only if the point $\mathcal{D}(h)$ lies above the hyperplane $\mathcal{D}(p)$.

5.2 H-Polytopes and V-Polytopes

A convex polytope in the plane is a convex polygon. Famous examples of convex polytopes in $\mathbb{R}^3$ are the Platonic solids: regular tetrahedron, cube, regular octahedron, regular dodecahedron, and regular icosahedron. A convex polytope in $\mathbb{R}^3$ is a convex set bounded by finitely many convex polygons. Such a set can be regarded as a convex hull of a finite point set, or as an intersection of finitely many half-spaces. We thus define two types of convex polytopes, based on these two views.

5.2.1 Definition (H-polytope and V-polytope). *An H-*polyhedron *is an intersection of finitely many closed half-spaces in some $\mathbb{R}^d$. An H-*polytope *is a bounded H-polyhedron.*

*A V-*polytope *is the convex hull of a finite point set in $\mathbb{R}^d$.*

A basic theorem about convex polytopes claims that from the mathematical point of view, H-polytopes and V-polytopes are equivalent.

5.2.2 Theorem. *Each V-polytope is an H-polytope. Each H-polytope is a V-polytope.*

This is one of the theorems that may look "obvious" and whose proof needs no particularly clever idea but does require some work. In the present case, we do not intend to avoid it. Actually, we have quite a neat proof in store, but we postpone it to the end of this section.

Although H-polytopes and V-polytopes are mathematically equivalent, there is an enormous difference between them from the computational point of view. That is, it matters a lot whether a convex polytope is given to us as a convex hull of a finite set or as an intersection of half-spaces. For example, given a set of n points specifying a V-polytope, how do we find its representation as an H-polytope? It is not hard to come up with some algorithm, but the problem is to find an efficient algorithm that would allow one to handle large real-world problems. This algorithmic question is not yet satisfactorily solved. Moreover, in some cases the number of required half-spaces may be astronomically large compared to the number n of points, as we will see later in this chapter.

As another illustration of the computational difference between V-polytopes and H-polytopes, we consider the maximization of a given linear function over a given polytope. For V-polytopes it is a trivial problem, since it suffices to substitute all points of V into the given linear function and select the maximum of the resulting values. But maximizing a linear function over the intersection of a collection of half-spaces is the basic problem of linear programming, and it is certainly nontrivial.

Terminology. The usual terminology does not distinguish V-polytopes and H-polytopes. A *convex polytope* means a point set in $\mathbb{R}^d$ that is a V-polytope (and thus also an H-polytope). An arbitrary, possibly unbounded, H-polyhedron is called a *convex polyhedron*. All polytopes

and polyhedra considered in this chapter are convex, and so the adjective "convex" is often omitted.

The *dimension* of a convex polyhedron is the dimension of its affine hull. It is the smallest dimension of a Euclidean space containing a congruent copy of P.

Basic examples. One of the easiest classes of polytopes is that of *cubes*. The d-dimensional cube as a point set is the Cartesian product $[-1,1]^d$.

$d=1$ $d=2$ $d=3$

As a V-polytope, the d-dimensional cube is the convex hull of the set $\{-1,1\}^d$ (2^d points), and as an H-polytope, it can be described by the inequalities $-1 \le x_i \le 1$, $i=1,2,\dots,d$, i.e., by $2d$ half-spaces. We note that it is also the unit ball of the maximum norm $\|x\|_\infty = \max_i |x_i|$.

Another important example is the class of *crosspolytopes* (or generalized octahedra). The d-dimensional crosspolytope is the convex hull of the "coordinate cross," i.e., $\mathrm{conv}\{e_1,-e_1,e_2,-e_2,\dots,e_d,-e_d\}$, where $e_1,\dots,e_d$ are the vectors of the standard orthonormal basis.

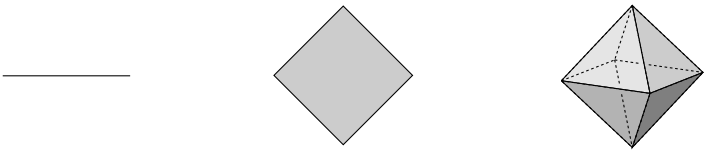

$d=1$ $d=2$ $d=3$

It is also the unit ball of the ℓ_1-norm $\|x\|_1 = \sum_{i=1}^d |x_i|$. As an H-polytope, it can be expressed by the 2^d half-spaces of the form $\langle \sigma, x\rangle \le 1$, where σ runs through all vectors in $\{-1,1\}^d$.

The polytopes with the smallest possible number of vertices (for a given dimension) are called simplices.

5.2.3 Definition (Simplex). *A* simplex *is the convex hull of an affinely independent point set in some* $\mathbb{R}^d$.

A d-dimensional simplex in $\mathbb{R}^d$ can also be represented as an intersection of $d+1$ half-spaces, as is not difficult to check.

A *regular* d-dimensional simplex is the convex hull of $d+1$ points with all pairs of points having equal distances.

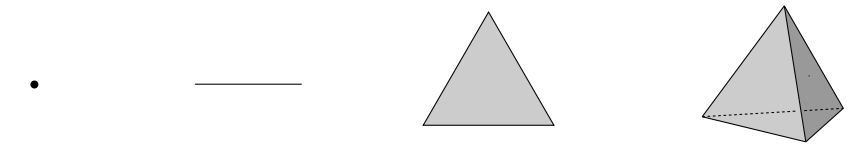

$d = 0$ $d = 1$ $d = 2$ $d = 3$

Unlike cubes and crosspolytopes, d-dimensional regular simplices do not have a very nice coordinate representation in $\mathbb{R}^d$. The simplest and most useful representation lives one dimension higher: The convex hull of the $d+1$ vectors $e_1, \ldots, e_{d+1}$ of the standard orthonormal basis in $\mathbb{R}^{d+1}$ is a d-dimensional regular simplex with side length $\sqrt{2}$.

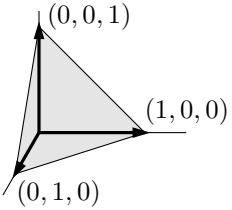

Proof of Theorem 5.2.2 (equivalence of H-polytopes and V-polytopes). We first show that any H-polytope is also a V-polytope. We proceed by induction on d. The case $d = 1$ being trivial, we suppose that $d \geq 2$.

So let Γ be a finite collection of closed half-spaces in $\mathbb{R}^d$ such that $P = \bigcap \Gamma$ is nonempty and bounded. For each $\gamma \in \Gamma$, let $F_\gamma = P \cap \partial\gamma$ be the intersection of P with the bounding hyperplane of γ. Each nonempty

F_γ is an H-polytope of dimension at most $d-1$ (correct?), and so it is the convex hull of a finite set $V_\gamma \subset F_\gamma$ by the inductive hypothesis.

We claim that $P = \mathrm{conv}(V)$, where $V = \bigcup_{\gamma\in\Gamma} V_\gamma$. Let $x \in P$ and let ℓ be a line passing through x. The intersection $\ell \cap P$ is a segment; let y and z be its endpoints. There are $\alpha, \beta \in \Gamma$ such that $y \in F_\alpha$ and $z \in F_\beta$ (if y were not on the boundary of any $\gamma \in \Gamma$, we could continue along ℓ a little further within P).

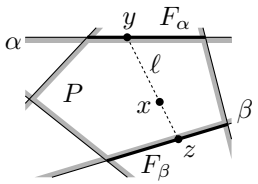

We have $y \in \mathrm{conv}(V_\alpha)$ and $z \in \mathrm{conv}(V_\beta)$, and thus $x \in \mathrm{conv}(V_\alpha \cup V_\beta) \subseteq \mathrm{conv}(V)$.

We have proved that any H-polytope is a V-polytope, and it remains to show that a V-polytope can be expressed as the intersection of finitely many half-spaces. This follows easily by duality (and implicitly uses the separation theorem), and we leave this as an exercise. □

5.3 Faces of a Convex Polytope

The surface of the 3-dimensional cube consists of 8 "corner" points called vertices, 12 edges, and 6 squares called *facets*. According to the perhaps more usual terminology in 3-dimensional geometry, the facets would be called faces. But in the theory of convex polytopes, the word face has a slightly different meaning, defined below. For the cube, not only the squares but also the vertices and the edges are all called *faces* of the cube.

5.3.1 Definition (Face). *A* face *of a convex polytope P is defined as*

- *either P itself, or*

- *a subset of P of the form $P \cap h$, where h is a hyperplane such that P is fully contained in one of the closed half-spaces determined by h.*

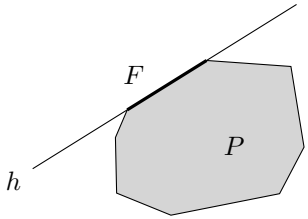

We observe that each face of P is a convex polytope. This is because P is the intersection of finitely many half-spaces and h is the intersection of two half-spaces, so the face is an H-polyhedron, and moreover, it is bounded.

If P is a polytope of dimension d, then its faces have dimensions -1, 0, 1, …, d, where -1 is, by definition, the dimension of the empty set. A face of dimension j is also called a j-face.

Names of faces. The 0-faces are called *vertices*, the 1-faces are called *edges*, and the $(d-1)$-faces of a d-dimensional polytope are called *facets*. The $(d-2)$-faces of a d-dimensional polytope are *ridges*; in the familiar 3-dimensional situation, edges = ridges. For example, the 3-dimensional cube has 28 faces in total: the empty face, 8 vertices, 12 edges, 6 facets, and the whole cube.

The following proposition shows that each V-polytope is the convex hull of its vertices, and that the faces can be described combinatorially: They are the convex hulls of certain subsets of vertices. This includes some intuitive facts such as that each edge connects two vertices.

A helpful notion is that of an *extremal point* of a set: For a set $X \subseteq \mathbb{R}^d$, a point $x \in X$ is extremal if $x \notin \operatorname{conv}(X \setminus \{x\})$.

5.3.2 Proposition. *Let $P \subset \mathbb{R}^d$ be a (bounded) convex polytope.*

(i) *("Vertices are extremal") The extremal points of P are exactly its vertices, and P is the convex hull of its vertices.*

(ii) *("Face of a face is a face") Let F be a face of P. The vertices of F are exactly those vertices of P that lie in F. More generally, the faces of F are exactly those faces of P that are contained in F.*

The proof is not essential for our further considerations, and it is given at the end of this section. below illustrates that things are not quite as simple as it might perhaps seem). The proposition has an appropriate analogue for polyhedra, but in order to avoid technicalities, we treat the bounded case only.

Graphs of polytopes. Each 1-dimensional face, or edge, of a convex polytope has exactly two vertices. We can thus define the *graph* $G(P)$ of a polytope P in the natural way: The vertices of the polytope are vertices of the graph, and two vertices are connected by an edge in the graph if they are vertices of the same edge of P. (The terms "vertices" and "edges" for graphs actually come from the corresponding notions for 3-dimensional convex polytopes.) Here is an example of a 3-dimensional polytope, the regular octahedron, with its graph:

For polytopes in $\mathbb{R}^3$, the graph is always planar: Project the polytope from its interior point onto a circumscribed sphere, and then make a "cartographic map" of this sphere, say by stereographic projection. Moreover, it can be shown that the graph is vertex 3-connected. (A graph G is called *vertex k-connected* if $|V(G)| \geq k+1$ and deleting any at most $k-1$ vertices leaves G connected.) Nicely enough, these properties characterize graphs of convex 3-polytopes:

5.3.3 Theorem (Steinitz theorem). *A finite graph is isomorphic to the graph of a 3-dimensional convex polytope if and only if it is planar and vertex 3-connected.*

We omit a proof of the considerably harder "if" part (exhibiting a polytope for every vertex 3-connected planar graph); all known proofs are quite complicated.

Graphs of higher-dimensional polytopes probably have no nice description comparable to the 3-dimensional case, and it is likely that the problem of deciding whether a given graph is isomorphic to a graph of a 4-dimensional convex polytope is NP-hard. It is known that the graph of every d-dimensional polytope is vertex d-connected (*Balinski's theorem*), but this is only a necessary condition.

Examples. A d-dimensional simplex has been defined as the convex hull of a $(d+1)$-point affinely independent set V. It is easy to see that each subset of V determines a face of the simplex. Thus, there are $\binom{d+1}{k+1}$ faces of dimension k, $k = -1, 0, \ldots, d$, and 2^{d+1} faces in total.

The d-dimensional crosspolytope has $V = \{e_1, -e_1, \ldots, e_d, -e_d\}$ as the vertex set. A proper subset $F \subset V$ determines a face if and only if there is no i such that both $e_i \in F$ and $-e_i \in F$ (exercise). It follows that there are 3^d+1 faces, including the empty one and the whole crosspolytope.

The nonempty faces of the d-dimensional cube $[-1, 1]^d$ correspond to vectors $v \in \{-1, 1, 0\}^d$. The face corresponding to such v has the vertex set $\{u \in \{-1, 1\}^d\colon u_i = v_i \text{ for all } i \text{ with } v_i \neq 0\}$. Geometrically, the vector v is the center of gravity of its face.

The face lattice. Let $\mathcal{F}(P)$ be the set of all faces of a (bounded) convex polytope P (including the empty face $\emptyset$ of dimension -1). We consider the partial ordering of $\mathcal{F}(P)$ by inclusion.

5.3.4 Definition (Combinatorial equivalence). *Two convex polytopes P and Q are called* combinatorially equivalent *if $\mathcal{F}(P)$ and $\mathcal{F}(Q)$ are isomorphic as partially ordered sets.*

We are going to state some properties of the partially ordered set $\mathcal{F}(P)$ without proofs. These are not difficult and are omitted.

It turns out that $\mathcal{F}(P)$ is a lattice (a partially ordered set satisfying additional axioms). We recall that this means the following two conditions:

- *Meets condition:* For any two faces $F, G \in \mathcal{F}(P)$, there exists a face $M \in \mathcal{F}(P)$, called the *meet* of F and G, that is contained in both F and G and contains all other faces contained in both F and G.

- *Joins condition:* For any two faces $F, G \in \mathcal{F}(P)$, there exists a face $J \in \mathcal{F}(P)$, called the *join* of F and G, that contains both F and G and is contained in all other faces containing both F and G.

The meet of two faces is their geometric intersection $F \cap G$.

For verifying the joins and meets conditions, it may be helpful to know that for a finite partially ordered set possessing the minimum element and the maximum element, the meets condition is equivalent to the joins condition, and so it is enough to check only one of the conditions.

Here is the face lattice of a 3-dimensional pyramid:

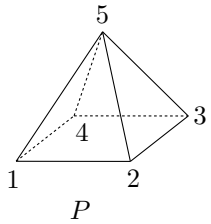

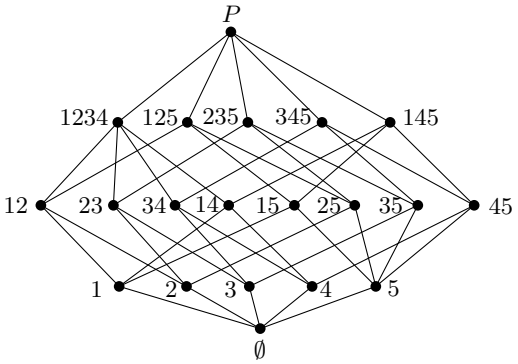

The vertices are numbered 1–5, and the faces are labeled by the vertex sets.

The face lattice is *graded*, meaning that every maximal chain has the same length (the rank of a face F is $\dim(F)+1$). Quite obviously, it is *atomic*: Every face is the join of its vertices. A little less obviously, it is *coatomic*; that is, every face is the meet (intersection) of the facets containing it. An important consequence is that combinatorial type of a polytope is determined by the vertex–facet incidences. More precisely, if we know the dimension and all subsets of vertices that are vertex sets of

facets (but without knowing the coordinates of the vertices, of course), we can uniquely reconstruct the whole face lattice in a simple and purely combinatorial way.

Face lattices of convex polytopes have several other nice properties, but no full algebraic characterization is known, and the problem of deciding whether a given lattice is a face lattice is algorithmically difficult (even for 4-dimensional polytopes).

The face lattice can be a suitable representation of a convex polytope in a computer. Each j-face is connected by pointers to its $(j-1)$-faces and to the $(j+1)$-faces containing it. On the other hand, it is a somewhat redundant representation: Recall that the vertex–facet incidences already contain the full information, and for some applications, even less data may be sufficient, say the graph of the polytope.

The dual polytope. Let P be a convex polytope containing the origin in its interior. Then the dual set P^* is also a polytope; we have verified this in the proof of Theorem 5.2.2.

5.3.5 Proposition. *For each $j = -1, 0, \ldots, d$, the j-faces of P are in a bijective correspondence with the $(d-j-1)$-faces of P^*. This correspondence also reverses inclusion; in particular, the face lattice of P^* arises by turning the face lattice of P upside down.*

Again we refer to the reader's diligence for a proof. Let us examine a few examples instead.

Among the five regular Platonic solids, the cube and the octahedron are dual to each other, the dodecahedron and the icosahedron are also dual, and the tetrahedron is dual to itself. More generally, if we have a 3-dimensional convex polytope and G is its graph, then the graph of the dual polytope is the dual graph to G, in the usual graph-theoretic sense. The dual of a d-simplex is a d-simplex, and the d-dimensional cube and the d-dimensional crosspolytope are dual to each other.

We conclude with two notions of polytopes "in general position."

5.3.6 Definition (Simple and simplicial polytopes). *A polytope P is called* simplicial *if each of its facets is a simplex (this happens, in particular, if the vertices of P are in general position, but general position is not necessary). A d-dimensional polytope P is called* simple *if each of its vertices is contained in exactly d facets.*

The faces of a simplex are again simplices, and so each proper face of a simplicial polytope is a simplex. Among the five Platonic solids, the tetrahedron, the octahedron, and the icosahedron are simplicial; and the tetrahedron, the cube, and the dodecahedron are simple. Crosspolytopes are simplicial, and cubes are simple. An example of a polytope that is neither simplicial nor simple is the 4-sided pyramid used in the illustration of the face lattice.

The dual of a simple polytope is simplicial, and vice versa. For a simple d-dimensional polytope, a small neighborhood of each vertex looks combinatorially like a neighborhood of a vertex of the d-dimensional cube. Thus, for each vertex v of a d-dimensional simple polytope, there are d edges emanating from v, and each k-tuple of these edges uniquely determines one k-face incident to v. Consequently, v belongs to $\binom{d}{k}$ k-faces, $k = 0, 1, \ldots, d$.

Proof of Proposition 5.3.2. In (i) ("vertices are extremal"), we assume that P is the convex hull of a finite point set. Among all such sets, we fix one that is inclusion-minimal and call it V_0. Let V_v be the vertex set of P, and let V_e be the set of all extremal points of P. We prove that $V_0 = V_v = V_e$, which gives (i). We have $V_e \subseteq V_0$ by the definition of an extremal point.

Next, we show that $V_v \subseteq V_e$. If $v \in V_v$ is a vertex of P, then there is a hyperplane h with $P \cap h = \{v\}$, and all of $P \setminus \{v\}$ lies in one of the open half-spaces defined by h. Hence $P \setminus \{v\}$ is convex, which means that v is an extremal point of P, and so $V_v \subseteq V_e$.

Finally we verify $V_0 \subseteq V_v$. Let $v \in V_0$; by the inclusion-minimality of V_0, we get that $v \notin C = \operatorname{conv}(V_0 \setminus \{v\})$. Since C and $\{v\}$ are disjoint compact convex sets, they can be strictly separated by a hyperplane h. Let h_v be the hyperplane parallel to h and containing v; this h_v has all

points of $V_0 \setminus \{v\}$ on one side.

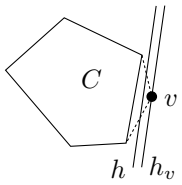

We want to show that $P \cap h_v = \{v\}$ (then v is a vertex of P and we are done). The set $P \setminus h_v = \mathrm{conv}(V_0) \setminus h_v$, being the intersection of a convex set with an open half-space, is convex. Any segment vx, where $x \in P \setminus h_v$, shares only the point v with the hyperplane h_v, and so $(P \setminus h_v) \cup \{v\}$ is convex as well. Since this set contains V_0 and is convex, it contains $P = \mathrm{conv}(V_0)$, and so $P \cap h_v = \{v\}$ indeed.

As for (ii) ("face of a face is a face"), it is clear that a face G of P contained in F is a face of F too (use the same witnessing hyperplane). For the reverse direction, we begin with the case of vertices. By a consideration similar to that at the end of the proof of (i), we see that $F = \mathrm{conv}(V) \cap h = \mathrm{conv}(V \cap h)$. Hence all the extremal points of F, which by (i) are exactly the vertices of F, are in V.

Finally, let F be a face of P defined by a hyperplane h, and let $G \subset F$ be a face of F defined by a hyperplane g within h; that is, g is a $(d-2)$-dimensional affine subspace of h with $G = g \cap F$ and with all of F on one side. Let γ be the closed half-space bounded by h with $P \subset \gamma$. We start rotating the boundary h of γ around g in the direction such that the rotated half-space γ' still contains F.

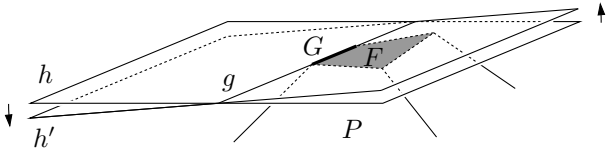

If we rotate by a sufficiently small amount, then all the vertices of P not lying in F are still in the interior of γ'. At the same time, the interior of γ' contains all the vertices of F not lying in G, while all the vertices of

G remain on the boundary h' of γ'. So h' defines a face of P (since all of P is on one side), and this face has the same vertex set as G, and so it equals G by the first part of (ii) proved above. □

5.4 Many Faces: The Cyclic Polytopes

A convex polytope P can be given to us by the list of vertices. How difficult is it to recover the full face lattice, or, more modestly, a representation of P as an intersection of half-spaces? The first question to ask is how large the face lattice or the collection of half-spaces can be, compared to the number of vertices. That is, what is the maximum total number of faces, or the maximum number of facets, of a convex polytope in $\mathbb{R}^d$ with n vertices? The dual question is, of course, the maximum number of faces or vertices of a bounded intersection of n half-spaces in $\mathbb{R}^d$.

Let $f_j = f_j(P)$ denote the number of j-faces of a polytope P. The vector $(f_0, f_1, \ldots, f_d)$ is called the *f-vector* of P. We thus assume $f_0 = n$ and we are interested in estimating the maximum value of f_{d-1} and of $\sum_{k=0}^{d} f_k$.

In dimensions 2 and 3, the situation is simple and favorable. For $d = 2$, our polytope is a convex polygon with n vertices and n edges, and so $f_0 = f_1 = n$, $f_2 = 1$. The f-vector is even determined uniquely.

A 3-dimensional polytope can be regarded as a drawing of a planar graph, in our case with n vertices. By well-known results for planar graphs, we have $f_1 \leq 3n-6$ and $f_2 \leq 2n-4$. Equalities hold if and only if the polytope is simplicial (all facets are triangles).

In both cases the total number of faces is linear in n. But as the dimension grows, polytopes become much more complicated. First of all, even the total number of faces of the most innocent convex polytope, the d-dimensional simplex, is *exponential* in d. But here we consider d fixed and relatively small, and we investigate the dependence on the number of vertices n.

Still, as we will see, for every $n \geq 5$ there is a 4-dimensional convex

polytope with n vertices and with every two vertices connected by an edge, i.e., with $\binom{n}{2}$ edges! This looks counterintuitive, but our intuition is based on the 3-dimensional case. In any fixed dimension d, the number of facets can be of order $n^{\lfloor d/2 \rfloor}$, which is rather disappointing for someone wishing to handle convex polytopes efficiently. On the other hand, complete desperation is perhaps not appropriate: Certainly not all polytopes exhibit this very bad behavior. For example, it is known that if we choose n points uniformly at random in the unit ball B^d, then the expected number of faces of their convex hull is only $o(n)$, for every fixed d.

It turns out that the number of faces for a given dimension and number of vertices is the largest possible for so-called *cyclic polytopes*, to be introduced next. First we define a very useful curve in $\mathbb{R}^d$.

5.4.1 Definition (Moment curve). *The curve* $\gamma = \{(t, t^2, \ldots, t^d)\colon t \in \mathbb{R}\}$ *in* $\mathbb{R}^d$ *is called the* moment curve.

5.4.2 Lemma. *Any hyperplane h intersects the moment curve γ in at most d points. If there are d intersections, then h cannot be tangent to γ, and thus at each intersection, γ passes from one side of h to the other.*

Proof. A hyperplane h can be expressed by the equation $\langle a, x\rangle = b$, or in coordinates $a_1x_1 + a_2x_2 + \cdots + a_dx_d = b$. A point of γ has the form $(t, t^2, \ldots, t^d)$, and if it lies in h, we obtain $a_1t + a_2t^2 + \cdots + a_dt^d - b = 0$. This means that t is a root of a nonzero polynomial $p_h(t)$ of degree at most d, and hence the number of intersections of h with γ is at most d. If there are d distinct roots, then they must be all simple. At a simple root, the polynomial $p_h(t)$ changes sign, and this means that the curve γ passes from one side of h to the other. □

As a corollary, we see that every $d+1$ points of the moment curve are affinely independent, for otherwise, we could pass a hyperplane through them. So the moment curve readily supplies *explicit* examples of point sets in general position.

5.4.3 Definition (Cyclic polytope). *The convex hull of finitely many points on the moment curve is called a* cyclic polytope.

How many facets does a cyclic polytope have? Each facet is determined by a d-tuple of vertices, and distinct d-tuples determine distinct facets. Here is a criterion telling us exactly which d-tuples determine facets.

5.4.4 Proposition (Gale's evenness criterion). *Let V be the vertex set of a cyclic polytope P considered with the linear ordering $\leq$ along the moment curve (larger vertices have larger values of the parameter t). Let $F = \{v_1, v_2, \ldots, v_d\} \subseteq V$ be a d-tuple of vertices of P, where $v_1 < v_2 < \cdots < v_d$. Then F determines a facet of P if and only if for any two vertices $u, v \in V \setminus F$, the number of vertices $v_i \in F$ with $u < v_i < v$ is even.*

Proof. Let h_F be the hyperplane affinely spanned by F. Then F determines a facet if and only if all the points of $V \setminus F$ lie on the same side of h_F.

Since the moment curve γ intersects h_F in exactly d points, namely at the points of F, it is partitioned into $d+1$ pieces, say $\gamma_0, \ldots, \gamma_d$, each lying completely in one of the half-spaces, as is indicated in the drawing:

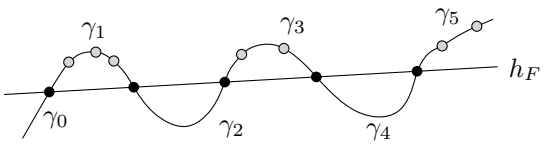

Hence, if the vertices of $V \setminus F$ are all contained in the odd-numbered pieces $\gamma_1, \gamma_3, \ldots$, as in the picture, or if they are all contained in the even-numbered pieces $\gamma_0, \gamma_2, \ldots$, then F determines a facet. This condition is equivalent to Gale's criterion. □

Now we can count the facets.

5.4.5 Theorem. *The number of facets of a d-dimensional cyclic polytope with n vertices ($n \geq d+1$) is*

$$\binom{n - \lfloor d/2 \rfloor}{\lfloor d/2 \rfloor} + \binom{n - \lfloor d/2 \rfloor - 1}{\lfloor d/2 \rfloor - 1} \text{ for } d \text{ even, and}$$

$$2\binom{n - \lfloor d/2 \rfloor - 1}{\lfloor d/2 \rfloor} \text{ for } d \text{ odd.}$$

For fixed d, this has the order of magnitude $n^{\lfloor d/2 \rfloor}$.

Proof. The number of facets equals the number of ways of placing d black circles and $n-d$ white circles in a row in such a way that we have an even number of black circles between each two white circles.

Let us say that an arrangement of black and white circles is *paired* if any contiguous segment of black circles has an even length (the arrangements permitted by Gale's criterion need not be paired because of the initial and final segments). The number of paired arrangements of $2k$ black circles and $n-2k$ white circles is $\binom{n-k}{k}$, since by deleting every second black circle we get a one-to-one correspondence with selections of the positions of k black circles among $n-k$ possible positions.

Let us return to the original problem, and first consider an *odd* $d = 2k+1$. In a valid arrangement of circles, we must have an odd number of consecutive black circles at the beginning or at the end (but not both). In the former case, we delete the initial black circle, and we get a paired arrangement of $2k$ black and $n-1-2k$ white circles. In the latter case, we similarly delete the black circle at the end and again get a paired arrangement as in the first case. This establishes the formula in the theorem for odd d.

For *even* $d = 2k$, the number of initial consecutive black circles is either odd or even. In the even case, we have a paired arrangement, which contributes $\binom{n-k}{k}$ possibilities. In the odd case, we also have an odd number of consecutive black circles at the end, and so by deleting the first and last black circles we obtain a paired arrangement of $2(k-1)$ black circles and $n-2k$ white circles. This contributes $\binom{n-k-1}{k-1}$ possibilities. □

5.5 The Upper Bound Theorem

The upper bound theorem, one of the earlier major achievements of the theory of convex polytopes, claims that the cyclic polytope has the largest possible number of faces.

5.5.1 Theorem (Upper bound theorem). *Among all d-dimensional convex polytopes with n vertices, the cyclic polytope maximizes the number of faces of each dimension.*

In this section we prove only an approximate result, which gives the correct order of magnitude for the maximum number of facets.

5.5.2 Proposition (Asymptotic upper bound theorem). *A d-dimensional convex polytope with n vertices has at most $2\binom{n}{\lfloor d/2\rfloor}$ facets and no more than $2^{d+1}\binom{n}{\lfloor d/2\rfloor}$ faces in total. For d fixed, both quantities thus have the order of magnitude $n^{\lfloor d/2\rfloor}$.*

First we establish this proposition for simplicial polytopes, in the following form.

5.5.3 Proposition. *Let P be a d-dimensional simplicial polytope. Then*

(a) $f_0(P) + f_1(P) + \cdots + f_d(P) \le 2^d f_{d-1}(P)$, *and*

(b) $f_{d-1}(P) \le 2f_{\lfloor d/2\rfloor -1}(P)$.

This implies Proposition 5.5.2 for simplicial polytopes, since the number of $(\lfloor d/2\rfloor -1)$-faces is certainly no bigger than $\binom{n}{\lfloor d/2\rfloor}$, the number of all $\lfloor d/2\rfloor$-tuples of vertices.

Proof of Proposition 5.5.3. We pass to the dual polytope P^*, which is simple. Now we need to prove $\sum_{k=0}^{d} f_k(P^*) \le 2^d f_0(P^*)$ and $f_0(P^*) \le 2f_{\lceil d/2\rceil}(P^*)$.

Each face of P^* has at least one vertex, and every vertex of a simple d-polytope is incident to 2^d faces, which gives the first inequality.

We now bound the number of vertices in terms of the number of $\lceil d/2 \rceil$-faces. This is the heart of the proof, and it shows where the mysterious exponent $\lfloor d/2 \rfloor$ comes from.

Let us rotate the polytope P^* so that no two vertices share the x_d-coordinate (i.e., no two vertices have the same vertical level).

Consider a vertex v with the d edges emanating from it. By the pigeonhole principle, there are at least $\lceil d/2 \rceil$ edges directed downwards or at least $\lceil d/2 \rceil$ edges directed upwards. In the former case, every $\lceil d/2 \rceil$-tuple of edges going up determines a $\lceil d/2 \rceil$-face for which v is the lowest vertex. In the latter case, every $\lceil d/2 \rceil$-tuple of edges going down determines a $\lceil d/2 \rceil$-face for which v is the highest vertex. Here is an illustration, unfortunately for the not too interesting 3-dimensional case, showing a situation with 2 edges going up and the corresponding 2-dimensional face having v as the lowest vertex:

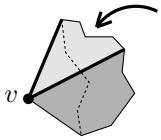

We have exhibited at least one $\lceil d/2 \rceil$-face for which v is the lowest vertex or the highest vertex. Since the lowest vertex and the highest vertex are unique for each face, the number of vertices is no more than twice the number of $\lceil d/2 \rceil$-faces. □

Warning. For simple polytopes, the total combinatorial complexity is proportional to the number of vertices, and for simplicial polytopes it is proportional to the number of facets (considering the dimension fixed, that is). For polytopes that are neither simple nor simplicial, the number of faces of intermediate dimensions can have larger order of magnitude than both the number of facets and the number of vertices.

Nonsimplicial polytopes. To prove the asymptotic upper bound theorem, it remains to deal with nonsimplicial polytopes. This is done by a perturbation argument, similar to numerous other results where general position is convenient for the proof but where we want to show that the result holds in degenerate cases as well.

5.5.4 Lemma. *For any d-dimensional convex polytope P there exists a d-dimensional simplicial polytope Q with $f_0(P) = f_0(Q)$ and $f_k(Q) \geq f_k(P)$ for all $k = 1, 2, \ldots, d$.*

Proof. The basic idea is very simple: Move (perturb) every vertex of P by a very small amount, in such a way that the vertices are in general position, and show that each k-face of P gives rise to at least one k-face of the perturbed polytope. There are several ways of doing this proof.

We process the vertices one by one. Let V be the vertex set of P and let $v \in V$. The operation of *ε-pushing* v is as follows: We choose a point v' lying in the interior of P, at distance at most ε from v, and on no hyperplane determined by the points of V, and we set $V' = (V \setminus \{v\}) \cup \{v'\}$. If we successively ε_v-push each vertex v of the polytope, the resulting vertex set is in general position and we have a simplicial polytope.

It remains to show that for any polytope P with vertex set V and any $v \in V$, there is an $\varepsilon > 0$ such that ε-pushing v does not decrease the number of faces.

Let $U \subset V$ be the vertex set of a k-face of P, $0 \leq k \leq d-1$, and let V' arise from V by ε-pushing v. If $v \notin U$, then no doubt, U determines a face of $\mathrm{conv}(V')$, and so we assume that $v \in U$. First suppose that v lies in the affine hull of $U \setminus \{v\}$; we claim that then $U \setminus \{v\}$ determines a k-face of $\mathrm{conv}(V')$. This follows from the following criterion, whose proof is left as an exercise: A subset $U \subset V$ is the vertex set of a face of $\mathrm{conv}(V)$ if and only if the affine hull of U is disjoint from $\mathrm{conv}(V \setminus U)$. We leave a detailed argument to the reader (one must use the fact that v is pushed inside).

If v lies outside of the affine hull of $U \setminus \{v\}$, then we want to show that $U' = (U \setminus \{v\}) \cup \{v'\}$ determines a k-face of $\mathrm{conv}(V')$. The affine hull of U is disjoint from the compact set $\mathrm{conv}(V \setminus U)$. If we move v continuously by a sufficiently small amount, the affine hull of U moves continuously, and so there is an $\varepsilon > 0$ such that if we move v within ε from its original position, the considered affine hull and $\mathrm{conv}(V \setminus U)$ remain disjoint. □

5.6 Voronoi Diagrams

Consider a finite set $P \subset \mathbb{R}^d$. For each point $p \in P$, we define a region $reg(p)$, which is the "sphere of influence" of the point p: It consists of the points $x \in \mathbb{R}^d$ for which p is the closest point among the points of P. Formally,

$$reg(p) = \{x \in \mathbb{R}^d\colon\ \mathrm{dist}(x,p) \le \mathrm{dist}(x,q) \text{ for all } q \in P\},$$

where $\mathrm{dist}(x,y)$ denotes the Euclidean distance of the points x and y. The *Voronoi diagram* of P is the set of all regions $reg(p)$ for $p \in P$. (More precisely, it is the cell complex induced by these regions; that is, every intersection of a subset of the regions is a face of the Voronoi diagram.) Here an example of the Voronoi diagram of a point set in the plane:

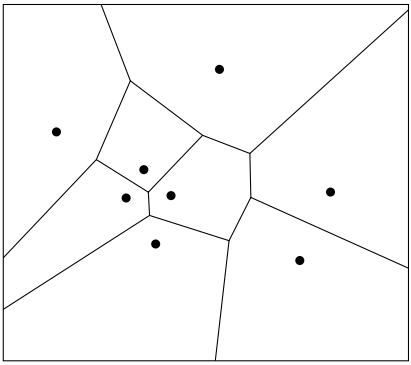

(Of course, the Voronoi diagram is clipped by a rectangle so that it fits into a finite page.) The points of P are traditionally called the *sites* in the context of Voronoi diagrams.

5.6.1 Observation. *Each region $reg(p)$ is a convex polyhedron with at most $|P|-1$ facets.*

Indeed,

$$reg(p) = \bigcap_{q \in P \setminus \{p\}} \{x\colon \mathrm{dist}(x,p) \le \mathrm{dist}(x,q)\}$$

is an intersection of $|P| - 1$ half-spaces. □

For $d = 2$, a Voronoi diagram of n points is a subdivision of the plane into n convex polygons (some of them are unbounded). It can be regarded as a drawing of a planar graph (with one vertex at the infinity, say), and hence it has a linear combinatorial complexity: n regions, $O(n)$ vertices, and $O(n)$ edges.

In the literature the Voronoi diagram also appears under various other names, such as the *Dirichlet tessellation*.

Examples of applications. Voronoi diagrams have been reinvented and used in various branches of science. Sometimes the connections are surprising. For instance, in archaeology, Voronoi diagrams help study cultural influences. Here we mention a few applications, mostly algorithmic.

- ("Post office problem" or nearest neighbor searching) Given a point set P in the plane, we want to construct a data structure that finds the point of P nearest to a given query point x as quickly as possible. This problem arises directly in some practical situations or, more significantly, as a subroutine in more complicated problems. The query can be answered by determining the region of the Voronoi diagram of P containing x. For this problem (point location in a subdivision of the plane), efficient data structures are known; see introductory texts on computational geometry.

- (Robot motion planning) Consider a disk-shaped robot in the plane. It should pass among a set P of point obstacles, getting from a given start position to a given target position and touching none of the obstacles.

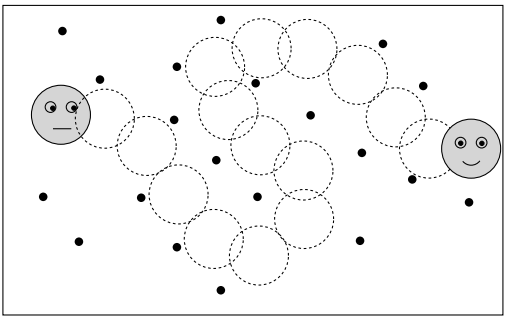

If such a passage is possible at all, the robot can always walk along the edges of the Voronoi diagram of P, except for the initial and final segments of the tour. This allows one to reduce the robot motion problem to a graph search problem: We define a subgraph of the Voronoi diagram consisting of the edges that are passable for the robot.

- (A nice triangulation: the Delaunay triangulation) Let $P \subset \mathbb{R}^2$ be a finite point set. In many applications one needs to construct a triangulation of P (that is, to subdivide $\text{conv}(P)$ into triangles with vertices at the points of P) in such a way that the triangles are not too skinny. Of course, for some sets, some skinny triangles are necessary, but we want to avoid them as much as possible. One particular triangulation that is usually very good, and provably optimal with respect to several natural criteria, is obtained as the dual graph to the Voronoi diagram of P. Two points of P are connected by an edge if and only if their Voronoi regions share an edge.

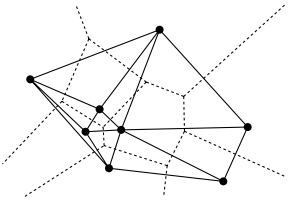

If no 4 points of P lie on a common circle then this indeed defines a triangulation, called the *Delaunay triangulation*[1] of P. The definition extends to points sets in $\mathbb{R}^d$ in a straightforward manner.

- (Interpolation) Suppose that $f: \mathbb{R}^2 \to \mathbb{R}$ is some smooth function whose values are known to us only at the points of a finite set $P \subset \mathbb{R}^2$. We would like to interpolate f over the whole polygon $\mathrm{conv}(P)$. Of course, we cannot really tell what f looks like outside P, but still we want a reasonable interpolation rule that provides a nice smooth function with the given values at P. Multidimensional interpolation is an extensive semiempirical discipline, which we do not seriously consider here; we explain only one elegant method based on Voronoi diagrams. To compute the interpolated value at a point $x \in \mathrm{conv}(P)$, we construct the Voronoi diagram of P, and we overlay it with the Voronoi diagram of $P \cup \{x\}$.

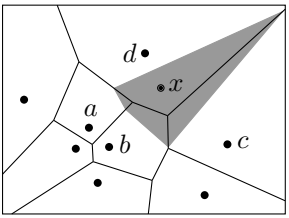

The region of the new point x cuts off portions of the regions of some of the old points. Let w_p be the area of the part of $reg(p)$ in the Voronoi diagram of P that belongs to $reg(x)$ after inserting x. The interpolated value $f(x)$ is

$$f(x) = \sum_{p \in P} \frac{w_p}{\sum_{q \in P} w_q} f(p).$$

An analogous method can be used in higher dimensions, too.

[1] Being a transcription from Russian, the spelling of Delaunay's name varies in the literature. For example, in crystallography literature he is usually spelled "Delone."

Relation of Voronoi diagrams to convex polyhedra. We now show that Voronoi diagrams in $\mathbb{R}^d$ correspond to certain convex polyhedra in $\mathbb{R}^{d+1}$.

First we define the *unit paraboloid* in $\mathbb{R}^{d+1}$:

$$U = \{x \in \mathbb{R}^{d+1}\colon x_{d+1} = x_1^2 + x_2^2 + \cdots + x_d^2\}.$$

For $d = 1$, U is a parabola in the plane.

In the sequel, let us imagine the space $\mathbb{R}^d$ as the hyperplane $x_{d+1} = 0$ in $\mathbb{R}^{d+1}$. For a point $p = (p_1, \ldots, p_d) \in \mathbb{R}^d$, let $\mathsf{e}(p)$ denote the hyperplane in $\mathbb{R}^{d+1}$ with equation

$$x_{d+1} = 2p_1x_1 + 2p_2x_2 + \cdots + 2p_dx_d - p_1^2 - p_2^2 - \cdots - p_d^2.$$

Geometrically, $\mathsf{e}(p)$ is the hyperplane tangent to the paraboloid U at the point $u(p) = (p_1, p_2, \ldots, p_d, p_1^2 + \cdots + p_d^2)$ lying vertically above p. It is perhaps easier to remember this geometric definition of $\mathsf{e}(p)$ and derive its equation by differentiation when needed. On the other hand, in the forthcoming proof we start out from the equation of $\mathsf{e}(p)$, and as a by-product, we will see that $\mathsf{e}(p)$ is the tangent to U at $u(p)$ as claimed.

5.6.2 Proposition. *Let $p, x \in \mathbb{R}^d$ be points and let $u(x)$ be the point of U vertically above x. Then $u(x)$ lies above the hyperplane $\mathsf{e}(p)$ or on it, and the vertical distance of $u(x)$ to $\mathsf{e}(p)$ is δ^2, where $\delta = \operatorname{dist}(x, p)$.*

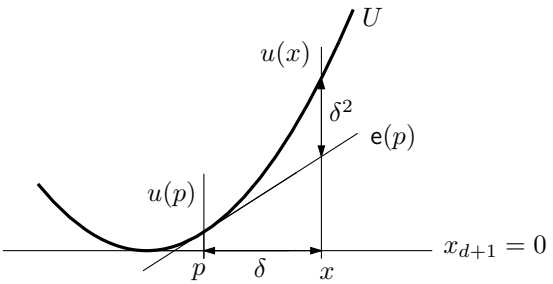

Proof. We just substitute into the equations of U and of $\mathsf{e}(p)$. The x_{d+1}-coordinate of $u(x)$ is $x_1^2 + \cdots + x_d^2$, while the x_{d+1}-coordinate of the point of $\mathsf{e}(p)$ above x is $2p_1x_1 + \cdots + 2p_dx_d - p_1^2 - \cdots - p_d^2$. The difference is $(x_1 - p_1)^2 + \cdots + (x_d - p_d)^2 = \delta^2$. □

Let $\mathcal{E}(p)$ denote the half-space lying above the hyperplane $\mathsf{e}(p)$. Consider an n-point set $P \subset \mathbb{R}^d$. By Proposition 5.6.2, $x \in reg(p)$ holds if and only if $\mathsf{e}(p)$ is vertically closest to U at x among all $\mathsf{e}(q)$, $q \in P$. Here is what we have derived:

5.6.3 Corollary. *The Voronoi diagram of P is the vertical projection of the facets of the polyhedron $\bigcap_{p\in P} \mathcal{E}(p)$ onto the hyperplane $x_{d+1} = 0$.* □

Here is an illustration for a planar Voronoi diagram:

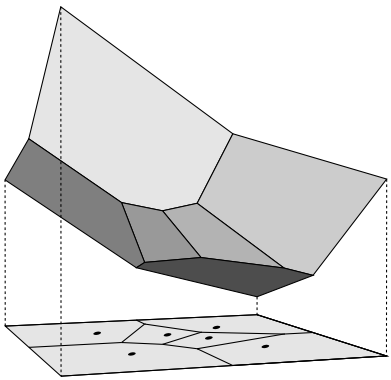

5.6.4 Corollary. *The maximum total number of faces of all regions of the Voronoi diagram of an n-point set in $\mathbb{R}^d$ is $O(n^{\lceil d/2 \rceil})$.*

Proof. We know that the combinatorial complexity of the Voronoi diagram equals the combinatorial complexity of an H-polyhedron with at most n facets in $\mathbb{R}^{d+1}$. By intersecting this H-polyhedron with a large simplex we can obtain a bounded polytope with at most $n+d+2$

facets, and we have not decreased the number of faces compared to the original H-polyhedron. Then the dual version of the asymptotic upper bound theorem (Theorem 5.5.2) implies that the total number of faces is $O(n^{\lceil d/2 \rceil})$, since $\lfloor (d+1)/2 \rfloor = \lceil d/2 \rceil$. □

The convex polyhedra in $\mathbb{R}^{d+1}$ obtained from Voronoi diagrams in $\mathbb{R}^d$ by the above construction are rather special, and so a lower bound for the combinatorial complexity of convex polytopes cannot be automatically transferred to Voronoi diagrams. But it turns out that the number of vertices of a Voronoi diagram on n points in $\mathbb{R}^d$ can really be of order $n^{\lceil d/2 \rceil}$.

The farthest-point Voronoi diagram. The projection of the H-polyhedron $\bigcap_{p \in P} \mathcal{E}(p)^{\mathrm{op}}$, where γ^{op} denotes the half-space opposite to γ, forms the *farthest-neighbor Voronoi diagram*, in which each point $p \in P$ is assigned the regions of points for which it is the farthest point. It can be shown that all nonempty regions of this diagram are unbounded and they correspond precisely to the points appearing on the surface of $\mathrm{conv}(P)$.

6

Number of Faces in Arrangements

Arrangements of lines in the plane and their higher-dimensional generalization, arrangements of hyperplanes in $\mathbb{R}^d$, are a basic geometric structure whose significance is comparable to that of convex polytopes. In fact, arrangements and convex polytopes are quite closely related: A cell in a hyperplane arrangement is a convex polyhedron, and conversely, each hyperplane arrangement in $\mathbb{R}^d$ corresponds canonically to a convex polytope in $\mathbb{R}^{d+1}$ of a special type, the so-called zonotope. But as is often the case with different representations of the same mathematical structure, convex polytopes and arrangements of hyperplanes emphasize different aspects of the structure and lead to different questions.

Whenever we have a problem involving a finite point set in $\mathbb{R}^d$ and partitions of the set by hyperplanes, we can use geometric duality, and we obtain a problem concerning a hyperplane arrangement. Arrangements appear in many other contexts as well; for example, some models of molecules give rise to arrangements of spheres in $\mathbb{R}^3$, and automatic planning of the motion of a robot among obstacles involves, implicitly or explicitly, arrangements of surfaces in higher-dimensional spaces.

Arrangements of hyperplanes have been investigated for a long time from various points of view. In several classical areas of mathematics one is mainly interested in topological and algebraic properties of the whole arrangement. Hyperplane arrangements are related to such marvelous

objects as Lie algebras, root systems, and Coxeter groups. In the theory of *oriented matroids* one studies the systems of sign vectors associated to hyperplane arrangements in an abstract axiomatic setting.

We are going to concentrate on estimating the combinatorial complexity (number of faces) in arrangements and neglect all the other directions.

General probabilistic techniques for bounding the complexity of geometric configurations constitute the second main theme of this chapter. These methods have been successful in attacking many more problems than can even be mentioned here.

6.1 Arrangements of Hyperplanes

Consider a finite set H of lines in the plane. They divide the plane into convex subsets of various dimensions, as is indicated in the following picture with 4 lines:

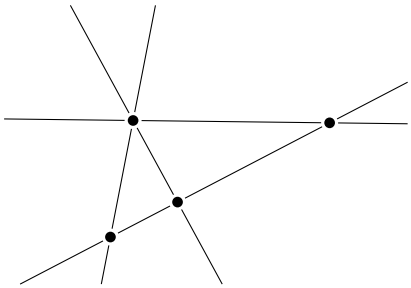

The intersections of the lines, indicated by black dots, are called the *vertices* or 0-faces. By removing all the vertices lying on a line $h \in H$, the line is split into two unbounded rays and several segments, and these parts are the *edges* or 1-faces. Finally, by deleting all the lines of H, the plane is divided into open convex polygons, called the *cells* or 2-faces. The vertices, edges, and cells are together called *faces*, and the arrangement of H is the collection of all these faces.[1]

[1] The terminology is not unified in the literature. What we call faces are sometimes referred to as cells (0-cells, 1-cells, and 2-cells).

An arrangement of a finite set H of hyperplanes in $\mathbb{R}^d$ is again a partition of $\mathbb{R}^d$ into relatively open convex faces. Their dimensions are 0 through d. As in the plane, the 0-faces are called vertices, the 1-faces edges, and the d-faces *cells*. Sometimes the $(d-1)$-faces are referred to as *facets*.

The cells are the connected components of $\mathbb{R}^d \setminus \bigcup H$. To obtain the facets, we consider the $(d-1)$-dimensional arrangements induced in the hyperplanes of H by their intersections with the other hyperplanes. That is, for each $h \in H$ we take the connected components of $h \setminus \bigcup_{h' \in H:\, h' \neq h} h'$. To obtain k-faces, we consider every possible k-flat L defined as the intersection of some $d-k$ hyperplanes of H. The k-faces of the arrangement lying within L are the connected components of $L \setminus \bigcup (H \setminus H_L)$, where $H_L = \{h \in H\colon L \subseteq h\}$.

Counting the cells in a hyperplane arrangement. We want to count the maximum number of faces in an arrangement of n hyperplanes in $\mathbb{R}^d$. As we will see, this is much simpler than the similar task for convex polytopes!

If a set H of hyperplanes is in general position, which means that the intersection of every k hyperplanes is $(d-k)$-dimensional, $k = 2, 3, \ldots, d+1$, the arrangement of H is called *simple*. For $|H| \geq d+1$ it suffices to require that every d hyperplanes intersect at a single point and no $d+1$ have a common point.

Every d-tuple of hyperplanes in a simple arrangement determines exactly one vertex, and so a simple arrangement of n hyperplanes has exactly $\binom{n}{d}$ vertices. We now calculate the number of cells; it turns out that the order of magnitude is also n^d for d fixed.

6.1.1 Proposition. *The number of cells (d-faces) in a simple arrangement of n hyperplanes in $\mathbb{R}^d$ equals*

$$\Phi_d(n) = \binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{d}. \tag{6.1}$$

First proof. We proceed by induction on the dimension d and the number of hyperplanes n. For $d = 1$ we have a line and n points in it.

These divide the line into $n+1$ one-dimensional pieces, and formula (6.1) holds. (The formula is also correct for $n = 0$ and all $d \geq 1$, since the whole space, with no hyperplanes, is a single cell.)

Now suppose that we are in dimension d, we have $n-1$ hyperplanes, and we insert another one. Since we assume general position, the $n-1$ previous hyperplanes divide the newly inserted hyperplane h into $\Phi_{d-1}(n-1)$ cells by the inductive hypothesis. Each such $(d-1)$-dimensional cell within h partitions one d-dimensional cell into exactly two new cells. The total increase in the number of cells caused by inserting h is thus $\Phi_{d-1}(n-1)$, and so

$$\Phi_d(n) = \Phi_d(n-1) + \Phi_{d-1}(n-1).$$

Together with the initial conditions (for $d = 1$ and for $n = 0$), this recurrence determines all values of Φ, and so it remains to check that formula (6.1) satisfies the recurrence. We have

$$\begin{aligned}\Phi_d(n-1) + \Phi_{d-1}(n-1) &= \tbinom{n-1}{0} + \left[\tbinom{n-1}{1} + \tbinom{n-1}{0}\right] \\ &\quad + \left[\tbinom{n-1}{2} + \tbinom{n-1}{1}\right] + \cdots + \left[\tbinom{n-1}{d} + \tbinom{n-1}{d-1}\right] \\ &= \tbinom{n-1}{0} + \tbinom{n}{1} + \tbinom{n}{2} + \cdots + \tbinom{n}{d} = \Phi_d(n).\end{aligned}$$

□

Second proof. This proof looks simpler, but a complete rigorous presentation is perhaps somewhat more demanding.

We proceed by induction on d, the case $d = 0$ being trivial. Let H be a set of n hyperplanes in $\mathbb{R}^d$ in general position; in particular, we assume that no hyperplane of H is horizontal and no two vertices of the arrangement have the same vertical level (x_d-coordinate).

Let g be an auxiliary horizontal hyperplane lying below all the vertices. A cell of the arrangement of H either is bounded from below, and in this case it has a unique lowest vertex, or is not bounded from below, and then it intersects g. The number of cells of the former type is the same as the number of vertices, which is $\binom{n}{d}$. The cells of the latter type correspond to the cells in the $(d-1)$-dimensional arrangement induced within g by the hyperplanes of H, and their number is thus $\Phi_{d-1}(n)$. □

What is the number of faces of the intermediate dimensions $1, 2, \ldots, d-1$ in a simple arrangement of n hyperplanes? This is not difficult to calculate using Proposition 6.1.1 (exercise); the main conclusion is that the total number of faces is $O(n^d)$ for a fixed d.

What about nonsimple arrangements? It turns out that a simple arrangement of n hyperplanes maximizes the number of faces of each dimension among arrangements of n hyperplanes. This can be verified by a perturbation argument, which is considerably simpler than the one for convex polytopes (Lemma 5.5.4), and which we omit.

6.2 Arrangements of Other Geometric Objects

Arrangements can be defined not only for hyperplanes but also for other geometric objects. For example, what is the arrangement of a finite set H of segments in the plane? As in the case of lines, it is a decomposition of the plane into faces of dimension $0, 1, 2$: the vertices, the edges, and the cells, respectively. The vertices are the intersections of the segments, the edges are the portions of the segments after removing the vertices, and the cells (2-faces) are the connected components of $\mathbb{R}^2 \setminus \bigcup H$. (Note that the endpoints of the segments are *not* included among the vertices.) While the cells of line arrangements are convex polygons, those in arrangements of segments can be complicated regions, even with holes:

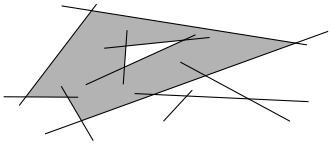

It is almost obvious that the total number of faces of the arrangement of n segments is at most $O(n^2)$. What is the maximum number of edges on the boundary of a single cell in such an arrangements? This seemingly innocuous question is surprisingly difficult (and we will not discuss it here).

Let us now present the definition of the arrangement for arbitrary sets $A_1, A_2, \ldots, A_n \subseteq \mathbb{R}^d$. The arrangement is a subdivision of space

into connected pieces again called the *faces*. Each face is an inclusion-maximal connected set that "crosses no boundary." More precisely, first we define an equivalence relation $\approx$ on $\mathbb{R}^d$: We put $x \approx y$ whenever x and y lie in the same subcollection of the A_i, that is, whenever $\{i\colon x \in A_i\} = \{i\colon y \in A_i\}$. So for each $I \subseteq \{1, 2, \ldots, n\}$, we have one possible equivalence class, namely $\{x \in \mathbb{R}^d\colon x \in A_i \Leftrightarrow i \in I\}$ (this is like a field in the Venn diagram of the A_i). But in typical geometric situations, most of the classes are empty. The faces of the arrangement of the A_i are the connected components of the equivalence classes. The reader is invited to check that for both hyperplane arrangements and arrangements of segments this definition coincides with the earlier ones.

Arrangements of algebraic surfaces. Quite often one needs to consider arrangements of the zero sets of polynomials. Let $p_1(x_1, x_2, \ldots, x_d), \ldots, p_n(x_1, x_2, \ldots, x_d)$ be polynomials with real coefficients in d variables, and let $Z_i = \{x \in \mathbb{R}^d\colon p_i(x) = 0\}$ be the zero set of p_i. Let D denote the maximum of the degrees of the p_i; when speaking of the arrangement of $Z_1, \ldots, Z_n$, one usually assumes that D is bounded by some (small) constant. Without a bound on D, even a single Z_i can have arbitrarily many connected components.

In many cases, the Z_i are algebraic surfaces, such as ellipsoids, paraboloids, etc., but since we are in the real domain, sometimes they need not look like surfaces at all. For example, the zero set of the polynomial $p(x_1, x_2) = x_1^2 + x_2^2$ consists of the single point $(0, 0)$. Although it is sometimes convenient to think of the Z_i as surfaces, the results stated below apply to zero sets of arbitrary polynomials of bounded degree.

It is known that if both d and D are considered as constants, the maximum number of faces in the arrangement of $Z_1, Z_2, \ldots, Z_n$ as above is at most $O(n^d)$. This is one of the most useful results about arrangements, with many surprising applications (a few are outlined below). In the literature one often finds a (formally weaker) version dealing with *sign patterns* of the polynomials p_i. A vector $\sigma \in \{-1, 0, +1\}^n$ is called a sign pattern of $p_1, p_2, \ldots, p_n$ if there exists an $x \in \mathbb{R}^d$ such that the sign of $p_i(x)$ is σ_i, for all $i = 1, 2, \ldots, n$. Trivially, the number of sign patterns for any n polynomials is at most 3^n. For $d = 1$, it is easy to

see that the actual number of sign patterns is much smaller, namely at most $2nD + 1$. It is not so easy to prove, but still true, that there are at most $C(d, D) \cdot n^d$ sign patterns in dimension d. This result is generally called the *Milnor–Thom theorem* (and it was apparently first proved by Oleinik and Petrovskiĭ, which fits the usual pattern in the history of mathematics). Here is a more precise (and more recent) version of this result, where the dependence on D and d is specified quite precisely.

6.2.1 Theorem (Number of sign patterns). *Let $p_1, p_2, \ldots, p_n$ be d-variate real polynomials of degree at most D. The number of faces in the arrangement of their zero sets $Z_1, Z_2, \ldots, Z_n \subseteq \mathbb{R}^d$, and consequently the number of sign patterns of $p_1, \ldots, p_n$ as well is at most $2(2D)^d \sum_{i=0}^{d} 2^i \binom{4n+1}{i}$. For $n \geq d \geq 2$, this expression is bounded by*

$$\left(\frac{50Dn}{d}\right)^d.$$

Proofs of these results are not included here because they would require at least one more chapter. They belong to the field of *real algebraic geometry*. The classical, deep, and extremely extensive field of *algebraic geometry* mostly studies algebraic varieties over algebraically closed fields, such as the complex numbers (and the questions of combinatorial complexity in our sense are not among its main interests). Real algebraic geometry investigates algebraic varieties and related concepts over the real numbers or other real-closed fields; the presence of ordering and the missing roots of polynomials make its flavor distinctly different.

Arrangements of pseudolines. An arrangement of pseudolines is a natural generalization of an arrangement of lines. Lines are replaced by curves, but we insist that these curves behave, in a suitable sense, like lines: For example, no two of them intersect more than once. This kind of generalization is quite different from, say, arrangements of planar algebraic curves, and so it perhaps does not quite belong to the present section. But besides mentioning pseudoline arrangements as a useful and interesting concept, we also need them for a (typical) example of

application of Theorem 6.2.1, and so we kill two birds with one stone by discussing them here.

An *(affine) arrangement of pseudolines* can be defined as the arrangement of a finite collection of curves in the plane that satisfy the following conditions:

(i) Each curve is x-monotone and unbounded in both directions; in other words, it intersects each vertical line in exactly one point.

(ii) Every two of the curves intersect in exactly one point and they cross at the intersection. (We do not permit "parallel" pseudolines, for they would complicate the definition unnecessarily.)[2]

The curves are called *pseudolines*, but while "being a line" is an absolute notion, "being a pseudoline" makes sense only with respect to a given collection of curves.

Here is an example of a (simple) arrangement of 5 pseudolines:

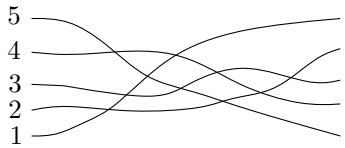

Much of what we have proved for arrangements of lines is true for arrangements of pseudolines as well. This holds for the maximum number of vertices, edges, and cells, but also for more sophisticated results like the Szemerédi–Trotter theorem on the maximum number of incidences

[2]This "affine" definition is a little artificial, and we use it only because we do not want to assume the reader's familiarity with the topology of the projective plane. In the literature one usually considers arrangements of pseudolines in the projective plane, where the definition is very natural: Each pseudoline is a closed curve whose removal does not disconnect the projective plane, and every two pseudolines intersect exactly once (which already implies that they cross at the intersection point). Moreover, one often adds the condition that the curves do not form a single *pencil*; i.e., not all of them have a common point, since otherwise, one would have to exclude the case of a pencil in the formulation of many theorems. But here we are not going to study pseudoline arrangements in any depth.

of m points and n lines; these results have proofs that do not use any properties of straight lines not shared by pseudolines.

One might be tempted to say that pseudolines are curves that behave topologically like lines, but as we will see below, in at least one sense this is profoundly wrong. The correct statement is that every two of them behave topologically like two lines, but arrangements of pseudolines are more general than arrangements of lines.

We should first point out that there is no problem with the "local" structure of the pseudolines, since each pseudoline arrangement can be redrawn equivalently (in a sense defined precisely below) by polygonal lines, as a *wiring diagram*:

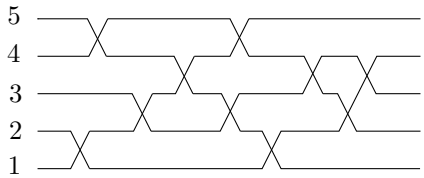

The difference between pseudoline arrangements and line arrangements is of a more global nature.

The arrangement of 5 pseudolines drawn above can be realized by straight lines:

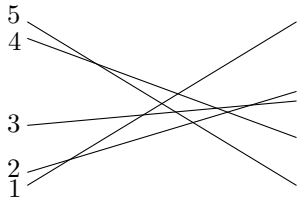

What is the meaning of "realization by straight lines"? To this end, we need a suitable notion of equivalence of two arrangements of pseudolines. There are several technically different possibilities; we again use an "affine" notion, one that is very simple to state but not the most common. Let H be a collection of n pseudolines. We number the pseudolines $1, 2, \ldots, n$ in the order in which they appear on the left of the

arrangement, say from the bottom to the top. For each i, we write down the numbers of the other pseudolines in the order they are encountered along the pseudoline i from left to right. For a simple arrangement we obtain a permutation π_i of $\{1, 2, \ldots, n\} \setminus \{i\}$ for each i. For the arrangement in the pictures, we have $\pi_1 = (2, 3, 5, 4)$, $\pi_2 = (1, 5, 4, 3)$, $\pi_3 = (1, 5, 4, 2)$, $\pi_4 = (5, 1, 3, 2)$, and $\pi_5 = (4, 1, 3, 2)$. For a nonsimple arrangement, some of the π_i are linear *quasiorderings*, meaning that several consecutive numbers can be chunked together. We call two arrangements *affinely isomorphic* if they yield the same $\pi_1, \ldots, \pi_n$, i.e., if each pseudoline meets the others in the same (quasi)order as the corresponding pseudoline in the other arrangement. Two affinely isomorphic pseudoline arrangements can be converted one to another by a suitable homeomorphism of the plane.

An arrangement of pseudolines is *stretchable* if it is affinely isomorphic to an arrangement of straight lines. It turns out that all arrangements of 8 or fewer pseudolines are stretchable, but there exists a nonstretchable arrangement of 9 pseudolines:

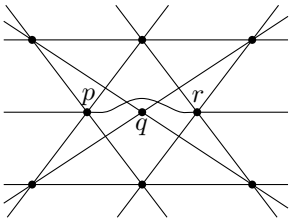

The proof of nonstretchability is based on the *Pappus theorem* in projective geometry, which states that if 8 straight lines intersect as in the drawing, then the points p, q, and r are collinear. By modifying this arrangement suitably, one can obtain a simple nonstretchable arrangement of 9 pseudolines as well.

Next, we show that most of the simple pseudoline arrangements are nonstretchable. The following construction shows that the number of isomorphism classes of simple arrangements of n pseudolines is at least $2^{\Omega(n^2)}$:

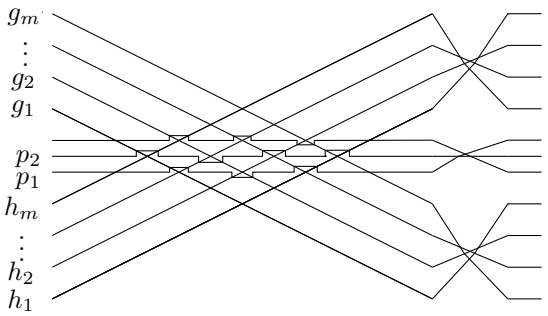

We have $m \approx \frac{n}{3}$, and the lines $h_1, \ldots, h_m$ and $g_1, \ldots, g_m$ form a regular grid. Each of the about $\frac{n}{3}$ pseudolines p_i in the middle passes near $\Omega(n)$ vertices of this grid, and for each such vertex it has a choice of going below it or above. This gives $2^{\Omega(n^2)}$ possibilities in total.

Now we use Theorem 6.2.1 to estimate the number of nonisomorphic simple arrangements of n straight lines. Let the lines be $\ell_1, \ldots, \ell_n$, where ℓ_i has the equation $y = a_i x + b_i$ and $a_1 > a_2 > \cdots > a_n$. The x-coordinate of the intersection $\ell_i \cap \ell_j$ is $\frac{b_i - b_j}{a_j - a_i}$. To determine the ordering π_i of the intersections along ℓ_i, it suffices to know the ordering of the x-coordinates of these intersections, and this can be inferred from the signs of the polynomials $p_{ijk}(a_i, b_i, a_j, b_j, a_k, b_k) = (b_i - b_j)(a_k - a_i) - (b_i - b_k)(a_j - a_i)$. So the number of nonisomorphic arrangements of n lines is no larger than the number of possible sign patterns of the $O(n^3)$ polynomials p_{ijk} in the $2n$ variables $a_1, b_1, \ldots, a_n, b_n$, and Theorem 6.2.1 yields the upper bound of $2^{O(n \log n)}$. For large n, this is a negligible fraction of the total number of simple pseudoline arrangements. (Similar considerations apply to nonsimple arrangements as well.)

The problem of deciding the stretchability of a given pseudoline arrangement has been shown to be algorithmically difficult (at least NP-hard). One can easily encounter this problem when thinking about line arrangements and drawing pictures: What we draw by hand are really pseudolines, not lines, and even with the help of a ruler it may be almost impossible to decide experimentally whether a given arrangement can really be drawn with straight lines. But there are computational methods that can decide stretchability in reasonable time at least for moderate

numbers of lines.

6.3 Number of Vertices of Level at Most k

In this section and the next one we investigate the maximum number of faces in certain naturally defined portions of hyperplane arrangements. We consider only simple arrangements, and we omit the (usually routine) perturbation arguments showing that simple arrangements maximize the investigated quantity.

Let H be a finite set of hyperplanes in $\mathbb{R}^d$, and assume that none of them is vertical, i.e., parallel to the x_d-axis. The *level* of a point $x \in \mathbb{R}^d$ is the number of hyperplanes of H lying strictly below x (the hyperplanes passing through x, if any, are not counted).

We are interested in the maximum possible number of vertices of level at most k in a simple arrangement of n hyperplanes. The following drawing shows the region of all points of level at most 2 in an arrangement of lines; we want to count the vertices lying in the region or on its boundary.

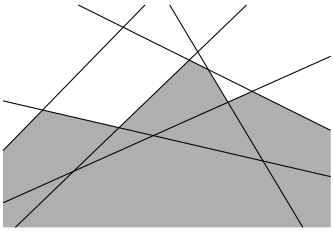

The vertices of level 0 are the vertices of the cell lying below all the hyperplanes, and since this cell is the intersection of at most n half-spaces, it has at most $O(n^{\lfloor d/2 \rfloor})$ vertices, by the asymptotic upper bound theorem (Theorem 5.5.2). From this result we derive a bound on the maximum number of vertices of level at most k. The elegant probabilistic technique used in the proof is generally applicable and probably more important than the particular result itself.

6.3.1 Theorem (Clarkson's theorem on levels). *The total number of vertices of level at most k in an arrangement of n hyperplanes in $\mathbb{R}^d$ is at most*

$$O(n^{\lfloor d/2 \rfloor}(k+1)^{\lceil d/2 \rceil}),$$

with the constant of proportionality depending on d.

We are going to prove the theorem for *simple* arrangements only. The general case can be derived from the result for simple arrangements by a standard perturbation argument. But let us stress that the simplicity of the arrangement is essential for the forthcoming proof.

For all k ($0 \le k \le n-d$), the bound is tight in the worst case. To see this for $k \ge 1$, consider a set of $\frac{n}{k}$ hyperplanes such that the lower unbounded cell in their arrangement is a convex polyhedron with $\Omega((\frac{n}{k})^{\lfloor d/2 \rfloor})$ vertices, and replace each of the hyperplanes by k very close parallel hyperplanes. Then each vertex of level 0 in the original arrangement gives rise to $\Omega(k^d)$ vertices of level at most k in the new arrangement.

A much more challenging problem is to estimate the maximum possible number of vertices of level *exactly* k (not discussed here).

One of the main motivations that led to Clarkson's theorem on levels was an algorithmic problem. Given an n-point set $P \subset \mathbb{R}^d$, we want to construct a data structure for fast answering of queries of the following type: For a query point $x \in \mathbb{R}^d$ and an integer t, report the t points of P that lie nearest to x.

Clarkson's theorem on levels is needed for bounding the maximum amount of memory used by a certain efficient algorithm. The connection is not entirely simple. It uses the lifting transform described in Section 5.6, relating the algorithmic problem in $\mathbb{R}^d$ to the complexity of levels in $\mathbb{R}^{d+1}$, and we do not discuss it here.

Proof of Theorem 6.3.1 for $d = 2$. We prove only this special case, for which the calculations are somewhat simpler.

Let H be a set of n lines in general position in the plane. Let p denote a certain suitable number in the interval $(0, 1)$ whose value will be determined at the end of the proof. Let us imagine the following random experiment. We choose a subset $R \subseteq H$ at random, by including

each line $h \in H$ into R with probability p, the choices being independent for distinct lines h.

Let us consider the arrangement of R, temporarily discarding all the other lines, and let $f(R)$ denote the number of vertices of level 0 in the arrangement of R. Since R is random, f is a random variable. We estimate the expectation of f, denoted by $\mathbf{E}[f]$, in two ways.

First, we have $f(R) \le |R|$ for any specific set R, and hence $\mathbf{E}[f] \le \mathbf{E}[|R|] = pn$.

Now we estimate $\mathbf{E}[f]$ differently: We bound it from below using the number of vertices of the arrangement of H of level at most k. For each vertex v of the arrangement of H, we define an event A_v meaning "v becomes one of the vertices of level 0 in the arrangement of R." That is, A_v occurs if v contributes 1 to the value of f. The event A_v occurs if and only if the following two conditions are satisfied:

- Both lines determining the vertex v lie in R.
- None of the lines of H lying below v falls into R.

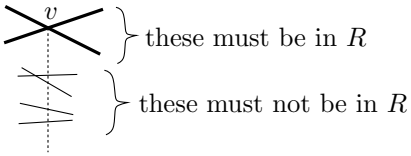

We deduce that $\mathrm{Prob}[A_v] = p^2(1-p)^{\ell(v)}$, where $\ell(v)$ denotes the level of the vertex v.

Let V be the set of all vertices of the arrangement of H, and let $V_{\le k} \subseteq V$ be the set of vertices of level at most k, whose cardinality we want to estimate. We have

$$\begin{aligned}\mathbf{E}[f] &= \sum_{v\in V} \mathrm{Prob}[A_v] \ge \sum_{v\in V_{\le k}} \mathrm{Prob}[A_v] \\ &= \sum_{v\in V_{\le k}} p^2(1-p)^{\ell(v)} \ge \sum_{v\in V_{\le k}} p^2(1-p)^k = |V_{\le k}| \cdot p^2(1-p)^k.\end{aligned}$$

Altogether we have derived $np \geq \mathbf{E}[f] \geq |V_{\leq k}| \cdot p^2(1-p)^k$, and so

$$|V_{\leq k}| \leq \frac{n}{p(1-p)^k}.$$

Let us now choose the number p so as to minimize the right-hand side. A convenient value is $p = \frac{1}{k+1}$; it does not yield the exact minimum, but it comes close. We have $\left(1 - \frac{1}{k+1}\right)^k \geq e^{-1} > \frac{1}{3}$ for all $k \geq 1$. This leads to $|V_{\leq k}| \leq 3(k+1)n$. □

Levels in arrangements. Besides vertices, we can consider all faces of level at most k, where the level of a face is the (common) level of all of its points. Using Theorem 6.3.1, it is not hard to prove that the number of all faces of level at most k in an arrangement of n hyperplanes is $O(n^{\lfloor d/2 \rfloor}(k+1)^{\lceil d/2 \rceil})$.

In the literature one often speaks about the *level* k in an arrangement of hyperplanes, meaning the boundary of the region of all points of level at most k. This is a polyhedral surface and each vertical line intersects it in exactly one point. It is a subcomplex of the arrangement; note that it may also contain faces of level different from k.

6.4 The Zone Theorem

Let H be a set of n hyperplanes in $\mathbb{R}^d$, and let g be a hyperplane that may or may not lie in H. The *zone* of g is the set of the faces of the arrangement of H that can see g. Here we imagine that the hyperplanes of H are opaque, and so we say that a face F can see the hyperplane g if there are points $x \in F$ and $y \in g$ such that the open segment xy is not intersected by any hyperplane of H (the face F is considered relatively open). Let us note that it does not matter which point $x \in F$ we choose: Either all of them can see g or none can. The picture shows the zone in a line arrangement:

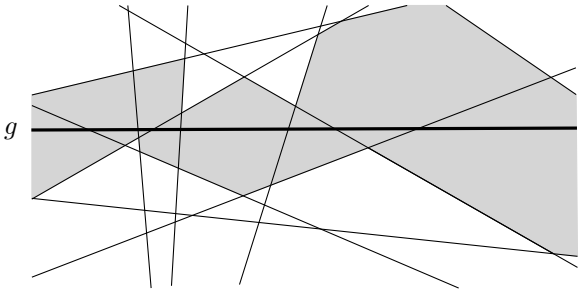

The following result bounds the maximum complexity of the zone. In the proof we will meet another interesting random sampling technique.

6.4.1 Theorem (Zone theorem). *The number of faces in the zone of any hyperplane in an arrangement of n hyperplanes in $\mathbb{R}^d$ is $O(n^{d-1})$, with the constant of proportionality depending on d.*

We prove the result only for simple arrangements; the general case follows, as usual, by a perturbation argument. Let us also assume that $g \notin H$ and that $H \cup \{g\}$ is in general position.

It is clear that the zone has $O(n^{d-1})$ *cells*, because each $(d-1)$-dimensional cell of the $(d-1)$-dimensional arrangement within g is intersects only one d-dimensional cell of the zone. On the other hand, this information is not sufficient to conclude that the total number of vertices of these cells is $O(n^{d-1})$: For example, it turns out that n arbitrarily chosen cells in an arrangement of n lines in the plane can together have as many as $\Omega(n^{4/3})$ vertices.

Proof. We proceed by induction on the dimension d. The base case is $d = 2$; it requires a separate treatment and does not follow from the trivial case $d = 1$ by the inductive argument shown below.

The case $d = 2$. Let H be a set of n lines in the plane in general position. We consider the zone of a line g. Since a convex polygon has the same number of vertices and edges, it suffices to bound the total number of 1-faces (edges) visible from the line g.

Imagine g drawn horizontally. We count the number of visible edges lying above g. Among those, at most n intersect the line g, since each

line of H gives rise to at most one such edge. The others are disjoint from g.

Consider an edge uv disjoint from g and visible from a point of g. Let $h \in H$ be the line containing uv, and let a be the intersection of h with g:

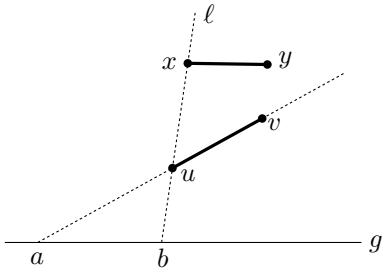

Let the notation be chosen in such a way that u is closer to a than v, and let $\ell \in H$ be the second line (besides h) defining the vertex u. Let b denote the intersection $\ell \cap g$. Let us call the edge uv a *right edge of the line* ℓ if the point b lies to the right of a, and a *left edge of the line* ℓ if b lies to the left of a.

We show that for each line ℓ there exists at most one right edge. If it were not the case, there would exist two edges, uv and xy, where u lies lower than x, which would both be right edges of ℓ, as in the above drawing. The edge xy should see some point of the line g, but the part of g lying to the right of a is obscured by the line h, and the part left of a is obscured by the line ℓ. This contradiction shows that the total number of right edges is at most n.

Symmetrically, we see that the number of left edges in the zone is at most n. The same bounds are obtained for edges of the zone lying below g. Altogether we have at most $O(n)$ edges in the zone, and the 2-dimensional case of the zone theorem is proved.

The case $d > 2$. Here we make the inductive step from $d-1$ to d. We assume that the total number of faces of a zone in $\mathbb{R}^{d-1}$ is $O(n^{d-2})$, and we want to bound the total number of zone faces in $\mathbb{R}^d$.

The first idea is to proceed by induction on n, bounding the maximum possible number of new faces created by adding a new hyperplane to

$n-1$ given ones. However, it is easy to find examples showing that the number of faces can increase roughly by n^{d-1}, and so this straightforward approach fails.

In the actual proof, we use a clever averaging argument. First, we demonstrate the method for the slightly simpler case of counting only the facets (i.e., $(d-1)$-faces) of the zone.

Let $f(n)$ denote the maximum possible number of $(d-1)$-faces in the zone in an arrangement of n hyperplanes in $\mathbb{R}^d$ (the dimension d is not shown in the notation in order to keep it simple). Let H be an arrangement and g a base hyperplane such that $f(n)$ is attained for them.

We consider the following random experiment. Color a randomly chosen hyperplane $h \in H$ red and the other hyperplanes of H blue. We investigate the expected number of *blue* facets of the zone, where a facet is blue if it lies in a blue hyperplane.

On the one hand, any facet has probability $\frac{n-1}{n}$ of becoming blue, and hence the expected number of blue facets is $\frac{n-1}{n} f(n)$.

We bound the expected number of blue facets in a different way. First, we consider the arrangement of blue hyperplanes only; it has at most $f(n-1)$ blue facets in the zone by the inductive hypothesis. Next, we add the red hyperplane, and we look by how much the number of blue facets in the zone can increase.

A new blue facet can arise by adding the red hyperplane only if the red hyperplane slices some existing blue facet F into two parts F_1 and F_2, as is indicated in the picture:

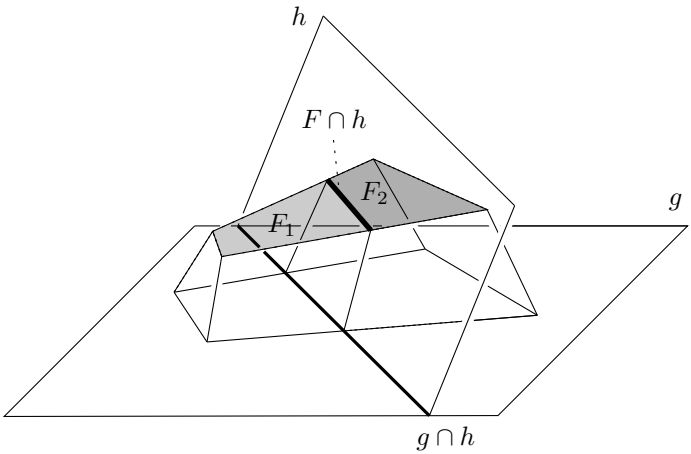

This increases the number of blue facets in the zone only if both F_1 and F_2 are visible from g. In such a case we look at the situation within the hyperplane h; we claim that $F \cap h$ is visible from $g \cap h$.

Let C be a cell of the zone in the arrangement of the blue hyperplanes having F on the boundary. We want to exhibit a segment connecting $F \cap h$ to $g \cap h$ within C. If $x_1 \in F_1$ sees a point $y_1 \in g$ and $x_2 \in F_2$ sees $y_2 \in g$, then the whole interior of the tetrahedron $x_1x_2y_1y_2$ is contained in C. The intersection of this tetrahedron with the hyperplane h contains a segment witnessing the visibility of $g \cap h$ from $F \cap h$.

If we intersect all the blue hyperplanes and the hyperplane g with the red hyperplane h, we get a $(d-1)$-dimensional arrangement, in which $F \cap h$ is a facet in the zone of the $(d-2)$-dimensional hyperplane $g \cap h$. By the inductive hypothesis, this zone has $O(n^{d-2})$ facets. Hence, adding h increases the number of blue facets of the zone by $O(n^{d-2})$, and so the total number of blue facets after h has been added is never more than $f(n-1) + O(n^{d-2})$.

We have derived the following inequality:

$$\frac{n-1}{n} f(n) \leq f(n-1) + O(n^{d-2}).$$

It implies $f(n) = O(n^{d-1})$, as we will demonstrate later for a slightly more general recurrence.

The previous considerations can be generalized for $(d-k)$-faces, where $1 \le k \le d-2$. Let $f_j(n)$ denote the maximum possible number of j-faces in the zone for n hyperplanes in dimension d. Let H be a collection of n hyperplanes where $f_{d-k}(n)$ is attained.

As before, we color one randomly chosen hyperplane $h \in H$ red and the others blue. A $(d-k)$-face is blue if its relative interior is disjoint from the red hyperplane. Then the probability of a fixed $(d-k)$-face being blue is $\frac{n-k}{n}$, and the expected number of blue $(d-k)$-faces in the zone is at most $\frac{n-k}{n} f_{d-k}(n)$.

On the other hand, we find that by adding the red hyperplane, the number of blue $(d-k)$-faces can increase by at most $O(n^{d-2})$, by the inductive hypothesis and by an argument similar to the case of facets. This yields the recurrence

$$\frac{n-k}{n} f_{d-k}(n) \le f_{d-k}(n-1) + O(n^{d-2}).$$

We use the substitution $\varphi(n) = \frac{f_{d-k}(n)}{n(n-1)\cdots(n-k+1)}$, which transforms our recurrence to $\varphi(n) \le \varphi(n-1) + O(n^{d-k-2})$. We assume $k < d-1$ (so the considered faces must not be edges or vertices). Then the last recurrence yields $\varphi(n) = O(n^{d-k-1})$, and hence $f_{d-k}(n) = O(n^{d-1})$.

For the case $k = d-1$ (edges), we would get only the bound $f_1(n) = O(n^{d-1} \log n)$ by this method. So the number of edges and vertices must be bounded by a separate argument, and we also have to argue separately for the planar case.

We are going to show that the number of vertices of the zone is at most proportional to the number of the 2-faces of the zone. Every vertex is contained in some 3-face of the zone. Within each such 3-face, the number of vertices is at most 3 times the number of 2-faces, because the 3-face is a 3-dimensional convex polyhedron. Since our arrangement is simple, each 2-face is contained in a bounded number of 3-faces. It follows that the total number of vertices is at most proportional to $f_2(n) = O(n^{d-1})$. The analogous bound for edges follows immediately from the bound for vertices. □