6

Number of Faces in Arrangements

Arrangements of lines in the plane and their higher-dimensional generalization, arrangements of hyperplanes in $\mathbf{R}^d$, 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 $\mathbf{R}^d$ corresponds canonically to a convex polytope in $\mathbf{R}^{d+1}$ 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 $\mathbf{R}^d$ 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 $\mathbf{R}^3$, 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 in this book. We begin with a simple but powerful sampling argument in Section 6.3 (somewhat resembling the proof of the crossing number theorem), add more tricks in Section 6.4, and finish with quite a sophisticated method, demonstrated on a construction of optimal $\frac{1}{r}$-cuttings, in Section 6.5.

6.1 Arrangements of Hyperplanes

We recall from Section 4.1 that for a finite set H of lines in the plane, the arrangement of H is a partition of the plane into relatively open convex subsets, the *faces* of the arrangement. In this particular case, the faces are the vertices (0-faces), the edges (1-faces), and the cells (2-faces).[1]

An arrangement of a finite set H of hyperplanes in $\mathbf{R}^d$ is again a partition of $\mathbf{R}^d$ 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 $\mathbf{R}^d \setminus \bigcup H$. 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 $H_L = \{h \in H \colon L \subseteq h\}$.

Remark on sign vectors. A face of the arrangement of H can be described by its *sign vector*. First we need to fix the *orientation* of each hyperplane $h \in H$. Each $h \in H$ partitions $\mathbf{R}^d$ into three regions: h itself and the two open half-spaces determined by it. We choose one of these open half-spaces as positive and denote it by $h^{\oplus}$, and we let the other one be negative, denoted by $h^{\ominus}$.

Let F be a face of the arrangement of H. We define the *sign vector* of F (with respect to the chosen orientations of the hyperplanes) as $\sigma(F) = (\sigma_h \colon h \in H)$, where

$$\sigma_h = \begin{cases} +1 & \text{if } F \subseteq h^{\oplus}, \\ 0 & \text{if } F \subseteq h, \\ -1 & \text{if } F \subseteq h^{\ominus}. \end{cases}$$

The sign vector determines the face F, since we have $F = \bigcap_{h \in H} h^{\sigma_h}$, where $h^0 = h$, $h^1 = h^{\oplus}$, and $h^{-1} = h^{\ominus}$. The following drawing shows the sign vectors of the marked faces in a line arrangement. Only the signs are shown, and the positive half-planes lie above their lines.

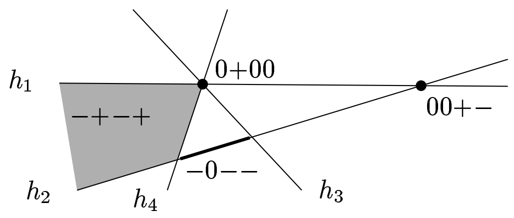

Of course, not all possible sign vectors correspond to nonempty faces. For n lines, there are 3^n sign vectors but only $O(n^2)$ faces, as we will derive below.

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

Counting the cells in a hyperplane arrangement. We want to count the maximum number of faces in an arrangement of n hyperplanes in $\mathbf{R}^d$. 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 $\mathbf{R}^d$ 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 $\mathbf{R}^d$ 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 1); 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.

Bibliography and remarks. The paper of Steiner [Ste26] from 1826 gives formulas for the number of faces in arrangements of lines, circles, planes, and spheres. Of course, his results have been extended in many ways since then (see, e.g., Zaslavsky [Zas75]). An early monograph on arrangements is Grünbaum [Grü72].

The questions considered in the subsequent sections, such as the combinatorial complexity of certain parts of arrangements, have been studied mainly in the last twenty years or so. A recent survey discussing a large part of the material of this chapter and providing many more facts and references is Agarwal and Sharir [AS00a].

The algebraic and topological investigation of hyperplane arrangements (both in real and complex spaces) is reflected in the book Orlik and Terao [OT91]. Let us remark that in these areas, one usually considers *central arrangements* of hyperplanes, where all the hyperplanes pass through the origin (and so they are linear subspaces of the underlying vector space). If such a central arrangement in $\mathbf{R}^d$ is intersected with a generic hyperplane not passing through the origin, one obtains a $(d-1)$-dimensional "affine" arrangement such as those considered by us. The correspondence is bijective, and so these two views of arrangements are not very different, but for many results, the formulation with central arrangements is more elegant.

The correspondence of arrangements to zonotopes is thoroughly explained in Ziegler [Zie94].

Exercises

1. (a) Count the number of faces of dimensions 1 and 2 for a simple arrangement of n planes in $\mathbf{R}^3$. $\boxed{2}$

   (b) Express the number of k-faces in a simple arrangement of n hyperplanes in $\mathbf{R}^d$. $\boxed{2}$

2. Prove that the number of *unbounded* cells in an arrangement of n hyperplanes in $\mathbf{R}^d$ is $O(n^{d-1})$ (for a fixed d). $\boxed{2}$

3. (a) Check that an arrangement of d or fewer hyperplanes in $\mathbf{R}^d$ has no bounded cell. $\boxed{2}$

   (b) Prove that an arrangement of $d+1$ hyperplanes in general position in $\mathbf{R}^d$ has exactly one bounded cell. $\boxed{3}$

4. How many d-dimensional cells are there in the arrangement of the $\binom{d}{2}$ hyperplanes in $\mathbf{R}^d$ with equations $\{x_i = x_j\}$, where $1 \le i < j \le d$? $\boxed{3}$

5. How many d-dimensional cells are there in the arrangement of the hyperplanes in $\mathbf{R}^d$ with the equations $\{x_i - x_j = 0\}$, $\{x_i - x_j = 1\}$, and $\{x_i - x_j = -1\}$, where $1 \le i < j \le d$? $\boxed{5}$

6. (Flags in arrangements)

   (a) Let H be a set of n lines in the plane, and let V be the set of vertices of their arrangement. Prove that the number of pairs (v, h) with $v \in V$, $h \in H$, and $v \in h$, i.e., the number of incidences $I(V, H)$, is bounded by $O(n^2)$. (Note that this is trivially true for simple arrangements.) $\boxed{2}$

   (b) Prove that the maximum number of d-tuples $(F_0, F_1, \ldots, F_d)$ in an arrangement of n hyperplanes in $\mathbf{R}^d$, where F_i is an i-dimensional face and F_{i-1} is contained in the closure of F_i, is $O(n^d)$ (d fixed). Such d-tuples are sometimes called *flags* of the arrangement. $\boxed{3}$

7. Let $P = \{p_1, \ldots, p_n\}$ be a point set in the plane. Let us say that points x, y have the *same view* of P if the points of P are visible in the same cyclic order from them. If rotating light rays emanate from x and from y, the points of P are lit in the same order by these rays. We assume that neither x nor y is in P and that neither of them can see two points of P in occlusion.

   (a) Show that the maximum possible number of points with mutually distinct views of P is $O(n^4)$. $\boxed{2}$

   (b) Show that the bound in (a) cannot be improved in general. $\boxed{4}$

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 $\mathbf{R}^2 \setminus \bigcup H$. (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:

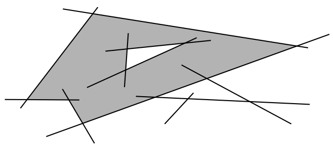

It is almost obvious that the total number of faces of the arrangement of n segments is at most $O(n^2)$. What is the maximum number of edges on the boundary of a single cell in such an arrangements? This seemingly innocuous question is surprisingly difficult, and most of Chapter 7 revolves around it.

Let us now present the definition of the arrangement for arbitrary sets $A_1, A_2, \ldots, A_n \subseteq \mathbf{R}^d$. 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 $\mathbf{R}^d$: 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 $\{x \in \mathbf{R}^d\colon x \in A_i \Leftrightarrow i \in I\}$ (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 $Z_i = \{x \in \mathbf{R}^d\colon p_i(x) = 0\}$ 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 and in the exercises). 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 $x \in \mathbf{R}^d$ 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 $2nD+1$ (Exercise 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 $Z_1, Z_2, \ldots, Z_n \subseteq \mathbf{R}^d$, 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:

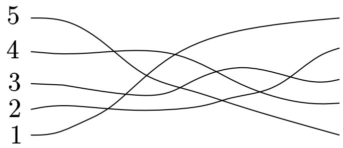

[2]This "affine" definition is a little artificial, and we use it only because we do not want to assume the reader's familiarity with the topology of the projective plane. In the literature one usually considers arrangements of pseudolines in the projective plane, where the definition is very natural: 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.

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 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*:

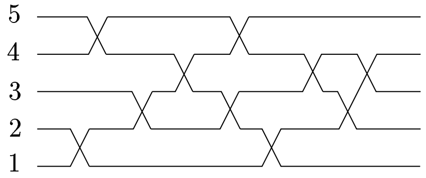

The difference between pseudoline arrangements and line arrangements is of a more global nature.

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

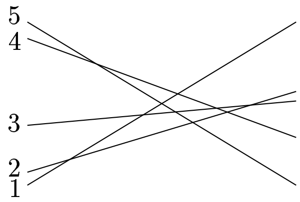

What is the meaning of "realization by straight lines"? To this end, we need a suitable notion of equivalence of two arrangements of pseudolines. There are several technically different possibilities; we again use an "affine" notion, one that is very simple to state but not the most common. Let H be 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.[3]

An arrangement of pseudolines is *stretchable* if it is affinely isomorphic to an arrangement of straight lines.[4] It turns out that all arrangements of 8 or fewer pseudolines are stretchable, but there exists a nonstretchable arrangement of 9 pseudolines:

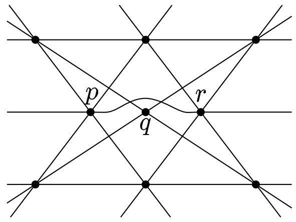

The proof of nonstretchability is based on the *Pappus theorem* in projective geometry, which states that if 8 straight lines intersect as in the drawing, then the points p, q, and r are collinear. By modifying this arrangement suitably, one can obtain a simple nonstretchable arrangement of 9 pseudolines 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 $2^{\Omega(n^2)}$:

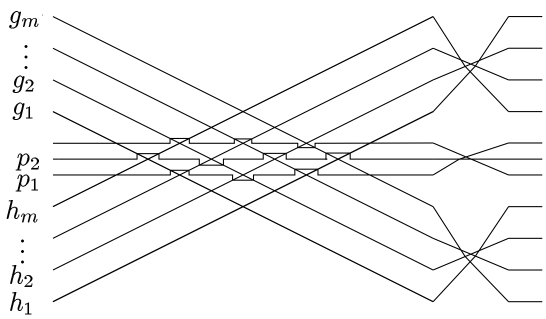

We have $m \approx \frac{n}{3}$, and the lines $h_1, \ldots, h_m$ and $g_1, \ldots, g_m$ form a regular grid. Each of the about $\frac{n}{3}$ pseudolines p_i in the middle passes near $\Omega(n)$ vertices of

[3]The more usual notion of *isomorphism* of pseudoline arrangements is defined for arrangements in the projective plane. The arrangement of H is isomorphic to the arrangement of H' if there exists a homeomorphism φ of the projective plane onto itself such that each pseudoline $h \in H$ is mapped to a pseudoline $\varphi(h) \in H'$. For affinely isomorphic arrangements in the affine plane, the corresponding arrangements in the projective plane are isomorphic, but the isomorphism in the projective plane also allows for mirror reflection and for "relocating the infinity." Combinatorially, the isomorphism in the projective plane can be described using the (quasi)orderings $\pi_1, \ldots, \pi_n$ as well. Here the π_i have to agree only up to a possible reversal and cyclic shift for each i, and also the numbering of the pseudolines by $1, 2, \ldots, n$ is not canonical.

We also remark that two arrangements of *lines* are isomorphic if and only if the dual point configurations have the same order type, up to a mirror reflection of the whole configuration (order types are discussed in Section 9.3).

[4]For isomorphism in the projective plane, one gets an equivalent notion of stretchability.

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.

Bibliography and remarks. A comprehensive account of real algebraic geometry is Bochnak, Coste, and Roy [BCR98]. Among the many available introductions to the "classical" algebraic geometry we mention the lively book Cox, Little, and O'Shea [CLO92].

The original bounds on the number of sign patterns, less precise than Theorem 6.2.1 but still implying the $O(n^d)$ bound for fixed d, were given independently by Oleinik and Petrovskiĭ [OP49], Milnor [Mil64], and Thom [Tho65]. Warren [War68] proved that the number of *d-dimensional cells* in the arrangement as in Theorem 6.2.1, and consequently the number of sign patterns consisting of ± 1's only, is at most $2(2D)^d \sum_{i=0}^{d} 2^i \binom{n}{i}$. The extension to faces of all dimensions, and to sign patterns including 0's, was obtained by Pollack and Roy [PR93].

Sometimes we have polynomials in many variables, but we are interested only in sign patterns attained at points that satisfy some additional algebraic conditions. Such a situation is covered by a result of Basu, Pollack, and Roy [BPR96]: The number of sign patterns attained by n polynomials of degree at most D on a k-dimensional algebraic variety $V \subseteq \mathbf{R}^d$, where V can be defined by polynomials of degree at most D, is at most $\binom{n}{k} O(D)^d$.

While bounding the number of sign patterns of multivariate polynomials appears complicated, there is a beautiful short proof of an almost tight bound on the number of *zero patterns*, due to Rónyai, Babai, and Ganapathy [RBG01], which we now sketch (in the simplest form, giving a slightly suboptimal result). A vector $\zeta \in \{0,1\}^n$ is a zero pattern of d-variate polynomials $p_1, \ldots, p_n$ with coefficients in a field F if there exists an $x = x(\zeta) \in F^d$ with $p_i(x) = 0$ exactly for the i with $\zeta_i = 0$. We show that if all the p_i have degree at most D, then the number of zero patterns cannot exceed $\binom{Dn+d}{d}$. For each zero pattern ζ, let q_ζ be the polynomial $\prod_{i:\, \zeta_i \neq 0} p_i$. We have $\deg q_\zeta \leq Dn$. Let us consider the q_ζ as elements of the vector space L of all d-variate polynomials over F of degree at most Dn. Using the basis of L consisting of all monomials of degree at most Dn, we obtain $\dim L \le \binom{Dn+d}{d}$. It remains to verify that the q_ζ are linearly independent (assuming that no p_i is identically 0). Suppose that $\sum_\zeta \alpha_\zeta q_\zeta = 0$ with $\alpha_\zeta \in F$ not all 0. Choose a zero pattern ξ with $\alpha_\xi \neq 0$ and with the largest possible number of 0's, and substitute $x(\xi)$ into $\sum_\zeta \alpha_\zeta q_\zeta$. This yields $\alpha_\xi = 0$, a contradiction.

Pseudoline arrangements. The founding paper is Levi [Lev26], where, among others, the nonstretchable arrangement of 9 lines drawn above was presented. A concise survey was written by Goodman [Goo97].

Pseudoline arrangements, besides being very natural, have also turned out to be a fruitful generalization of line arrangements. Some problems concerning line arrangements or point configurations were first solved only in the more general setting of pseudoline arrangements, and certain algorithms for line arrangements, the so-called *topological sweep* methods, use an auxiliary pseudoline to speed up the computation; see [Goo97].

Infinite families of pseudolines have been considered as well, and even *topological planes*, which are analogues of the projective plane but made of pseudolines. It is known that every finite configuration of pseudolines can be extended to a topological plane, and there are uncountably many distinct topological planes; see Goodman, Pollack, Wenger, and Zamfirescu [GPWZ94].

Oriented matroids. The possibility of representing each pseudoline arrangement by a wiring diagram makes it clear that a pseudoline arrangement can also be considered as a purely combinatorial object. The appropriate combinatorial counterpart of a pseudoline arrangement is called an oriented matroid of rank 3. More generally, similar to arrangements of pseudolines, one can define arrangements of *pseudohyperplanes* in $\mathbf{R}^d$, and these are combinatorially captured by oriented matroids of rank $d+1$. Here the rank is one higher than the space dimension, because an oriented matroid of rank d is usually viewed as a combinatorial abstraction of a *central* arrangement of hyperplanes in $\mathbf{R}^d$ (with all hyperplanes passing through 0).

There are several different but equivalent definitions of an oriented matroid. We present a definition in the so-called *covector form.* An *oriented matroid* is a set $\mathcal{V}^* \subseteq \{-1,0,1\}^n$ that is symmetric ($v \in \mathcal{V}^*$ implies $-v \in \mathcal{V}^*$), contains the zero vector, and satisfies the following two more complicated conditions:

- (Closed under composition) If $u, v \in \mathcal{V}^*$, then $u \circ v \in \mathcal{V}^*$, where $(u \circ v)_i = u_i$ if $u_i \neq 0$ and $(u \circ v)_i = v_i$ if $u_i = 0$.
- (Admits elimination) If $u, v \in \mathcal{V}^*$ and $j \in S(u,v) = \{i\colon u_i = -v_i \neq 0\}$, then there exists $w \in \mathcal{V}^*$ such that $w_j = 0$ and $w_i = (u \circ v)_i$ for all $i \notin S(u,v)$.

The *rank* of an oriented matroid $\mathcal{V}^*$ is the largest r such that there is an increasing chain $v_1 \prec v_2 \prec \cdots \prec v_r$, $v_i \in \mathcal{V}^*$, where $u \preceq v$ means $u_i \preceq v_i$ for all i and where $0 \prec 1$ and $0 \prec -1$. At first sight, all this may look quite mysterious, but it becomes much clearer if one thinks of a basic example, where $\mathcal{V}^*$ is the set of sign vectors of all faces of a central arrangement of hyperplanes in $\mathbf{R}^d$.

It turns out that every oriented matroid of rank 3 corresponds to an arrangement of pseudolines. More generally, *Lawrence's representation theorem* asserts that every oriented matroid of rank d comes from some central arrangement of pseudohyperplanes in $\mathbf{R}^d$, and so the purely combinatorial notion of oriented matroid corresponds, essentially uniquely, to the topological notion of a (central) arrangement of pseudohyperplanes.[5]

Oriented matroids are also naturally obtained from configurations of points or vectors. In the notation of Section 5.6 (Gale transform), if $\bar{\boldsymbol{a}}$ is a sequence of n

[5]The correspondence need not really be one-to-one. For example, the oriented matroids of two projectively isomorphic pseudoline arrangements agree only up to reorientation.

vectors in $\mathbf{R}^r$, then both the sets $\mathrm{sgn}(\mathsf{LinVal}(\bar{a}))$ and $\mathrm{sgn}(\mathsf{LinDep}(\bar{a}))$ are oriented matroids in the sense of the above definition. The first one has rank r, and the second, rank $n-r$.

We are not going to say much more about oriented matroids, referring to Ziegler [Zie94] for a quick introduction and to Björner, Las Vergnas, Sturmfels, White, and Ziegler [BVS+99] for a comprehensive account.

Stretchability. The following results illustrate the surprising difficulty of the stretchability problem for pseudoline arrangements. They are analogous to the statements about realizability of 4-dimensional convex polytopes mentioned in Section 5.3, and they were actually found much earlier.

Certain (simple) stretchable arrangements of n pseudolines require coefficients with $2^{\Omega(n)}$ digits in the equations of the lines, in every straight-line realization (Goodman, Pollack, and Sturmfels [GPS90]). Deciding the stretchability of a given pseudoline arrangement is NP-hard (Shor [Sho91] has a relatively simple proof), and in fact, it is polynomially equivalent to the problem of solvability of a system of polynomial inequalities with integer coefficients. This follows from results of Mnëv, published in Russian in 1985 (proofs were only sketched; see [Mne89] for an English version). This work went unnoticed in the West for some time, and so some of the results were rediscovered by other authors.

Although detailed proofs of such theorems are technically demanding, the principle is rather simple. Given two real numbers, suitably represented by geometric quantities, one can produce their sum and their product by classical geometric constructions by ruler. (Since ruler constructions are invariant under projective transformations, the numbers are represented as cross-ratios.) By composing such constructions, one can express the solvability of $p(x_1,\ldots,x_n)=0$, for a given n-variate polynomial p with integer coefficients, by the stretchability of a suitable arrangement in the projective plane. Dealing with inequalities and passing to simple arrangements is somewhat more complicated, but the idea is similar.

Practical algorithms for deciding stretchability have been studied extensively by Bokowski and Sturmfels [BS89] and by Richter-Gebert (see, e.g., [RG99]).

Mnëv [Mne89] was mainly interested in the *realization spaces* of arrangements. Let H be a fixed stretchable arrangement. Each straight-line arrangement H' affinely isomorphic to H can be represented by a point in $\mathbf{R}^{2n}$, with the $2n$ coordinates specifying the coefficients in the equations of the lines of H'. Considering all possible H' for a given H, we obtain a subset of $\mathbf{R}^{2n}$. For some time it was conjectured that this set, the realization space of H, has to be path-connected, which would mean that one straight-line realization could be converted to any other by a continuous motion while retaining the affine isomorphism type.[6] Not only is this false, but the realization space can have arbitrarily many components. In a suitable sense, it can even have arbitrary topological type. Whenever $A\subseteq\mathbf{R}^n$ is a set definable by a formula involving finitely many polynomial inequalities with integer coefficients, Boolean connectives, and quantifiers, there is a line arrangement whose realization space S is homotopy equivalent to A (Mnëv's main result actually talks about the stronger notion of *stable equivalence* of S and A; see, e.g., [Goo97] or [BVS+99]). Similar theorems were proved by Richter-Gebert for the realization spaces of 4-dimensional polytopes [RG99], [RG97].

These results for arrangements and polytopes can be regarded as instances of a vague but probably quite general principle: *"Almost none of the combinatorially*

[6] In fact, these questions have been studied mainly for the isomorphism of arrangements in the projective plane. There one has to be a little careful, since a mirror reflection can easily make the realization space disconnected, and so the mirror reflection (or the whole action of the general linear group) is factored out first.

imaginable geometric configurations are geometrically realizable, and it is difficult to decide which ones are." Of course, there are exceptions, such as the graphs of 3-dimensional convex polytopes.

Encoding pseudoline arrangements. The lower bound $2^{\Omega(n^2)}$ for the number of isomorphism classes of pseudoline arrangements is asymptotically tight. Felsner [Fel97] found a nice encoding of such an arrangement by an $n \times n$ matrix of 0's and 1's, from which the isomorphism type can be reconstructed: The entry (i, j) of the matrix is 1 iff the jth leftmost crossing along the pseudoline number i is with a pseudoline whose number k is larger than i.

Exercises

1. Let $p_1(x), \ldots, p_n(x)$ be univariate real polynomials of degree at most D. Check that the number of sign patterns of the p_i is at most $2nD+1$. $\boxed{2}$

2. (Intersection graphs) Let S be a set of n line segments in the plane. The *intersection graph* of S is the graph on n vertices, which correspond to the segments of S, with two vertices connected by an edge if and only if the corresponding two segments intersect.

   (a) Prove that the graph obtained from K_5 by subdividing each edge exactly once is not the intersection graph of segments in the plane (and not even the intersection graph of any arcwise connected sets in the plane). $\boxed{4}$

   (b) Use Theorem 6.2.1 to prove that *most* graphs are not intersection graphs of segments: While the total number of graphs on n given vertices is $2^{\binom{n}{2}} = 2^{n^2/2+O(n)}$, only $2^{O(n \log n)}$ of them are intersection graphs of segments (be careful about collinear segments!). $\boxed{3}$

   (c) Show that the number of (isomorphism classes of) intersection graphs of planar arcwise connected sets, and even of planar convex sets, on n vertices *cannot* be bounded by $2^{O(n \log n)}$. (The right order of magnitude does not seem to be known for either of these classes of intersection graphs.) $\boxed{4}$

3. (Number of combinatorially distinct simplicial convex polytopes) Use Theorem 6.2.1 to prove that for every dimension $d \geq 3$ there exists $C_d > 0$ such that the number of combinatorial types of *simplicial* polytopes in $\mathbf{R}^d$ with n vertices is at most $2^{C_d n \log n}$. (The combinatorial equivalence means isomorphic face lattices; see Definition 5.3.4.) $\boxed{4}$

   Such a result was proved by Alon [Alo86b] and by Goodman and Pollack [GP86].

4. (Sign patterns of matrices and rank) Let A be a real $n \times n$ matrix. The *sign matrix* $\sigma(A)$ is the $n \times n$ matrix with entries in $\{-1, 0, +1\}$ given by the signs of the corresponding entries in A.

   (a) Check that A has rank at most q if and only if there exist $n \times q$ matrices U and V with $A = UV^T$. $\boxed{3}$

   (b) Estimate the number of distinct sign matrices of rank q using Theorem 6.2.1, and conclude that there exists an $n \times n$ matrix S containing only entries $+1$ and -1 such that any real matrix A with $\sigma(A) = S$ has rank at least cn, with a suitable constant $c > 0$. $\boxed{3}$

   The result in (b) is from Alon, Frankl, and Rödl [AFR85] (for another application see [Mat96b]).

5. (Extendible pseudosegments) A *family of pseudosegments* is a finite collection $S = \{s_1, s_2, \ldots, s_n\}$ of curves in the plane such that each s_i is x-monotone and its vertical projection on the x-axis is a closed interval, every two curves in the

family intersect at most once, and whenever they intersect they cross (tangential contacts are not allowed). Such an S is called *extendible* if there is a family $L = \{\ell_1, \dots, \ell_n\}$ of pseudolines such that $s_i \subseteq \ell_i$, $i = 1, 2, \dots, n$.

(a) Find an example of a nonextendible family of 3 pseudosegments. $\boxed{1}$

(b) Define an oriented graph G with vertex set S and with an edge from s_i to s_j if $s_i \cap s_j \neq \emptyset$ and s_i is below s_j on the left of their intersection. Check that if S is extendible, then G is acyclic. $\boxed{1}$

(c) Prove that, conversely, if G is acyclic, then S is extendible. Extend the pseudosegments one by one, maintaining the acyclicity of G. $\boxed{3}$

(d) Let I_i be the projection of s_i on the x-axis. Show that if for every $i < j$, $I_i \cap I_j = \emptyset$ or $I_i \subseteq I_j$ or $I_j \subseteq I_i$, then G is acyclic, and hence S is extendible. $\boxed{2}$

(e) Given a family of closed intervals $I_1, \dots, I_n \subseteq \mathbf{R}$, show that each interval in the family can be partitioned into at most $O(\log n)$ subintervals in such a way that the resulting family of subintervals has the property as in (d). This implies that an arbitrary family of n pseudosegments can be cut into a family of $O(n \log n)$ extendible pseudosegments. $\boxed{3}$

These notions and results are from Chan [Cha00a].

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 $\mathbf{R}^d$, and assume that none of them is vertical, i.e., parallel to the x_d-axis. The *level* of a point $x \in \mathbf{R}^d$ is the number of hyperplanes of H lying strictly below x (the hyperplanes passing through x, if any, are not counted). This extends the definition for lines from Section 4.7.

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.

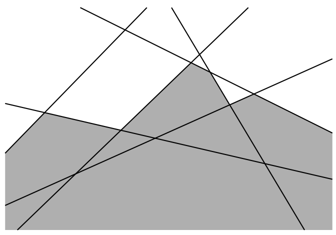

The vertices of level 0 are the vertices of the cell lying below all the hyperplanes, and since this cell is the intersection of at most n half-spaces, it has at most $O(n^{\lfloor d/2 \rfloor})$ vertices, by the asymptotic upper bound theorem (Theorem 5.5.2). From this result we derive a bound on the 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 $\mathbf{R}^d$ 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. This will be discussed in Chapter 11.

One of the main motivations that led to Clarkson's theorem on levels was an algorithmic problem. Given an n-point set $P \subset \mathbf{R}^d$, we want to construct a data structure for fast answering of queries of the following type: For a query point $x \in \mathbf{R}^d$ 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 5.7, relating the algorithmic problem in $\mathbf{R}^d$ to the complexity of levels in $\mathbf{R}^{d+1}$, and we do not discuss it here.

Proof of Theorem 6.3.1 for $d = 2$. First we demonstrate 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.

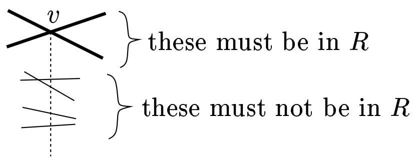

We deduce that $\mathrm{Prob}[A_v] = p^2(1-p)^{\ell(v)}$, where $\ell(v)$ denotes the level of the vertex v.

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

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

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 $|V_{\leq k}| \leq 3(k+1)n$. □

Proof for an arbitrary dimension. The idea of the proof is the same as above. As for the technical realization, there are at least two possible routes. The first is to retain the same probability distribution for selecting the sample R (picking each hyperplane of the given set H independently with probability p); in this case, most of the proof remains as before, but we need a lemma showing that $\mathbf{E}\left[|R|^{\lfloor d/2 \rfloor}\right] = O((pn)^{\lfloor d/2 \rfloor})$. This is not difficult to prove, either from a Chernoff-type inequality or by elementary calculations (see Exercises 6.5.2 and 6.5.3).

The second possibility, which we use here, is to change the probability distribution. Namely, we define an integer parameter r and choose a random r-element subset $R \subseteq H$, with all the $\binom{n}{r}$ subsets being equally probable.

With this new way of choosing R, we proceed as in the proof for $d = 2$. We define $f(R)$ as the number of vertices of level 0 in the arrangement of R and estimate $\mathbf{E}[f]$ in two ways. On the one hand, we have $f(R) = O(r^{\lfloor d/2 \rfloor})$ for all R, and so

$$\mathbf{E}[f] = O(r^{\lfloor d/2 \rfloor}).$$

The notation V for the set of all vertices of the arrangement of H, $V_{\leq k}$ for the vertices of level at most k, and A_v for the event "v is a vertex of level 0 in the arrangement of R," is as in the previous proof. The conditions for A_v are

- All the d hyperplanes defining the vertex v fall into R.
- None of the hyperplanes of H lying below v fall into R.

So if $\ell = \ell(v)$ is the level of v, then

$$\text{Prob}[A_v] = \frac{\binom{n-d-\ell}{r-d}}{\binom{n}{r}}.$$

For brevity, we denote this quantity by $P(\ell)$. We note that it is a *decreasing* function of ℓ. Therefore,

$$\mathbf{E}[f] = \sum_{v \in V} \text{Prob}[A_v] \geq |V_{\leq k}| \cdot P(k).$$

Combining with $\mathbf{E}[f] = O(r^{\lfloor d/2 \rfloor})$ derived earlier, we obtain

$$|V_{\leq k}| \leq \frac{O(r^{\lfloor d/2 \rfloor})}{P(k)}. \tag{6.2}$$

An appropriate value for the parameter r is $r = \lfloor \frac{n}{k+1} \rfloor$. (This is not surprising, since in the previous proof, the size of R was concentrated around $pn = \frac{n}{k+1}$.) Then we have the following estimate:

6.3.2 Lemma. *Suppose that* $1 \leq k \leq \frac{n}{2d} - 1$, *which implies* $2d \leq r \leq \frac{n}{2}$. *Then*

$$P(k) \geq c_d (k+1)^{-d}$$

for a suitable $c_d > 0$ *depending only on* d.

We postpone the proof of the lemma a little and finish the proof of Theorem 6.3.1. We want to substitute the bound from the lemma into (6.2). In order to meet the assumptions of the lemma, we must restrict the range of k somewhat. But if, say, $k \geq \frac{n}{2d}$, then the bound claimed by the theorem is of order n^d and thus trivial, and for $k = 0$ we already know that the theorem holds. So we may assume $1 \leq k \leq \frac{n}{2d} - 1$, and we have

$$|V_{\leq k}| \leq O(r^{\lfloor d/2 \rfloor})(k+1)^{-d} = O\left(n^{\lfloor d/2 \rfloor}(k+1)^{\lceil d/2 \rceil}\right).$$

This establishes the theorem. □

Proof of Lemma 6.3.2.

$$\begin{aligned}
P(k) &= \frac{\binom{n-d-k}{r-d}}{\binom{n}{r}} \\
&= \frac{(n-d-k)(n-d-k-1)\cdots(n-k-r+1)}{n(n-1)\cdots(n-r+1)} \cdot r(r-1)\cdots(r-d+1) \\
&= \frac{r(r-1)\cdots(r-d+1)}{n(n-1)\cdots(n-d+1)} \cdot \frac{n-d-k}{n-d} \cdot \frac{n-d-k-1}{n-d-1} \cdots \frac{n-k-r+1}{n-r+1} \\
&\geq \left(\frac{r}{2n}\right)^d \left(1 - \frac{k}{n-d}\right)\left(1 - \frac{k}{n-d-1}\right) \cdots \left(1 - \frac{k}{n-r+1}\right) \\
&\geq \left(\frac{r}{2n}\right)^d \left(1 - \frac{k}{n-r+1}\right)^r.
\end{aligned}$$

Now, $\frac{r}{n} \geq (\frac{n}{k+1} - 1)/n \geq \frac{1}{2(k+1)}$ (since $k < \frac{n}{2}$, say) and $1 - \frac{k}{n-r+1} \geq 1 - \frac{2k}{n}$ (a somewhat finer calculation actually gives $1 - \frac{k+1}{n}$ here). Since $k \leq \frac{n}{4}$, we can use the inequality $1-x \geq e^{-2x}$ valid for $x \in [0, \frac{1}{2}]$, and we arrive at

$$P(k) \geq (\tfrac{r}{dn})^d e^{-4kr/n} \geq c_d(k+1)^{-d}.$$

Lemma 6.3.2 is proved. □

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. In Section 4.7 we considered such levels in arrangements of lines.

Bibliography and remarks. Clarkson's theorem on levels was first proved in Clarkson [Cla88a] (see Clarkson and Shor [CS89] for the journal version). The elegant proof technique has many other applications, and we will meet it several more times, combined with additional tricks into sophisticated arguments. The theorem can be formulated in an abstract framework outlined in the notes to Section 6.5. New variations on the basic method were noted by Sharir [Sha01] (see Exercises 4 and 5).

In the planar case, the $O(nk)$ bound on the complexity of levels 0 through k was known before Clarkson's paper, apparently first proved by Goodman and Pollack [GP84]. Alon and Györi [AG86] determined the exact constant of proportionality (which Clarkson's proof in the present form cannot provide). Welzl [Wel01] proved an exact upper bound in $\mathbf{R}^3$; see the notes to Section 11.3 for a little more about his method. Several other related references can be found, e.g., in Agarwal and Sharir [AS00a].

Exercises

1. Show that for n hyperplanes in $\mathbf{R}^d$ in general position, the total number of vertices of levels $k, k+1, \ldots, n-d$ is at most $O(n^{\lfloor d/2 \rfloor}(n-k)^{\lceil d/2 \rceil})$. $\boxed{2}$

2. (a) Consider n lines in the plane in general position (their arrangement is simple). Call a vertex v of their arrangement an *extreme* if one of its defining lines has a positive slope and the other one has a negative slope. Prove that there are at most $O((k+1)^2)$ extremes of level at most k. Imitate the proof of Clarkson's theorem on levels. $\boxed{2}$

   (b) Show that the bound in (a) cannot be improved in general. $\boxed{1}$

3. Let $K_1, \ldots, K_n$ be circular disks in the plane. Show that the number of intersections of their boundary circles that are contained in at most k disks is bounded by $O(nk)$. Use the result of Exercise 5.7.10 and assume general position if convenient. $\boxed{3}$

4. Let L be a set of n nonvertical lines in the plane in general position.

(a) Let W be an arbitrary subset of vertices of the arrangement of L, and let X_W be the number of pairs (v, ℓ), where $v \in W$, $\ell \in L$, and ℓ goes (strictly) below v. For every real number $p \in (0,1)$, prove that $X_W \geq p^{-1}|W| - p^{-2}n$. $\boxed{3}$

(b) Let W be a set of vertices in the arrangement of L such that no line of L lies strictly below more than k vertices of W, where $k \geq 1$. Use (a) to prove $|W| = O(n\sqrt{k}\,)$. $\boxed{2}$

(c) Check that the bound in (b) is tight for all $k \leq \frac{n}{2}$. $\boxed{2}$

This exercise and the next one are from Sharir [Sha01].

5. Let P be an n-point set in the plane in general position (no 4 points on a common circle). Let C be a set of circles such that each circle in C passes through 3 points of P and contains no more than k points of P in its interior. Prove that $|C| \leq O(nk^{2/3})$, by an approach analogous to that of Exercise 4. $\boxed{3}$

6.4 The Zone Theorem

Let H be a set of n hyperplanes in $\mathbf{R}^d$, 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:

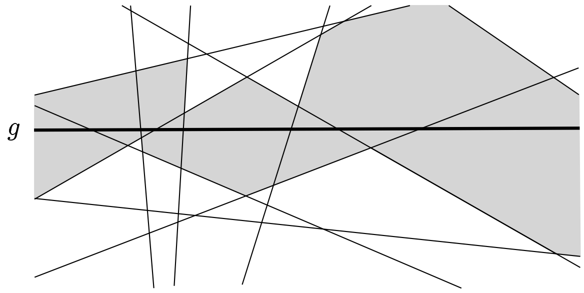

The following result bounds the maximum complexity of the zone. In the proof we will meet another interesting random sampling technique.

6.4.1 Theorem (Zone theorem). *The number of faces in the zone of any hyperplane in an arrangement of n hyperplanes in $\mathbf{R}^d$ 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, as we know from Chapter 4, 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 $d = 2$. (For another proof see Exercise 7.1.5.) 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:

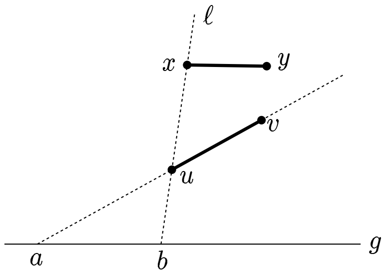

Let the notation be chosen in such a way that u is closer to a than v, and let $\ell \in H$ be the second line (besides h) defining the vertex u. Let b denote the intersection $\ell \cap g$. Let us call the edge uv a *right edge of the line* ℓ if the point b lies to the right of a, and a *left edge of 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 $\mathbf{R}^{d-1}$ is $O(n^{d-2})$, and we want to bound the total number of zone faces in $\mathbf{R}^d$.

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 $\mathbf{R}^d$ (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:

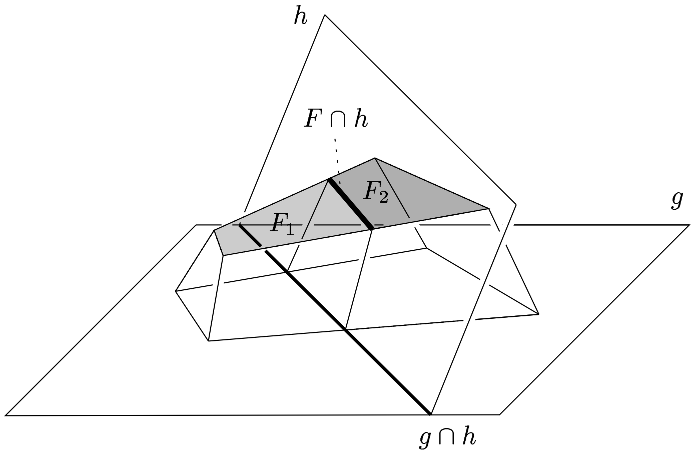

This increases the number of blue facets in the zone only if both F_1 and F_2 are visible from g. In such a case we look at the situation within the hyperplane h; we claim that $F \cap h$ is visible from $g \cap h$.

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 $1 \leq k \leq d-2$. 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

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

We use the substitution $\varphi(n) = \frac{f_{d-k}(n)}{n(n-1)\cdots(n-k+1)}$, which transforms our recurrence to $\varphi(n) \leq \varphi(n-1) + O(n^{d-k-2})$. 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. □

Zones in other arrangements. The maximum complexity of a zone can be investigated for objects other than hyperplanes. We can consider two classes $\mathcal{Z}$ and $\mathcal{A}$ of geometric objects in $\mathbf{R}^d$ and ask for the maximum complexity of the zone of a $\zeta \in \mathcal{Z}$ in the arrangement of n objects $a_1, a_2, \ldots, a_n \in \mathcal{A}$. This leads to a wide variety of problems. For some of them, interesting results have been obtained by extending the technique shown above.

Most notably, if ζ is a k-flat in $\mathbf{R}^d$, $0 \le k \le d$, or more generally, a *k-dimensional algebraic variety* in $\mathbf{R}^d$ of degree bounded by a constant, then the zone of ζ in an arrangement of n *hyperplanes* has complexity at most

$$O\left(n^{\lfloor (d+k)/2 \rfloor}(\log n)^\beta\right),$$

where $\beta = 1$ for $d + k$ odd and $\beta = 0$ for $d + k$ even. (The logarithmic factor seems likely to be superfluous in this bound; perhaps a more sophisticated proof could eliminate it.) With ζ being a k-flat, this result can be viewed as an interpolation between the asymptotic upper bound theorem and the zone theorem: For $k = 0$, with ζ being a single point, we consider the complexity of a single cell, while for $k = d-1$, we have the zone of a hyperplane. The key ideas of the proof are outlined in the notes below; for a full proof we refer to the literature.

A simple trick relates the zone problem to another question, the maximum complexity of a single cell in an arrangement. For example, what is the complexity of the zone of a segment ζ in an arrangement of n *line segments*? On the one hand, ζ can be chosen as a single point, and so the maximum zone complexity is at least the maximum possible complexity of a cell in an arrangement of n segments. On the other hand, the complexity of the zone of ζ is no more than the maximum cell complexity in an arrangement of $2n$ segments, since we can split each segment by making a tiny hole near the intersection with ζ:

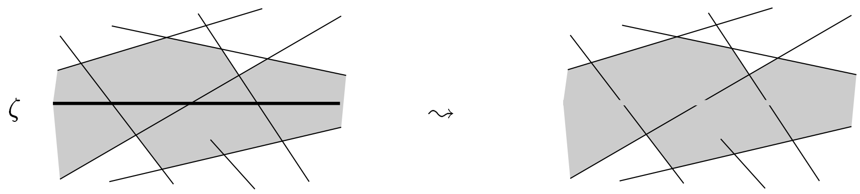

A similar reduction works for the zone of a triangle in an arrangement of triangles in $\mathbf{R}^3$ and in many other cases. Results presented in Section 7.6 will show that under quite general assumptions, the zone complexity in dimension d is no more than $O(n^{d-1+\varepsilon})$, for an arbitrarily small (but fixed) $\varepsilon > 0$.

Bibliography and remarks. The two-dimensional zone theorem was established by Chazelle, Guibas, and Lee [CGL85], with the proof shown above, and independently by Edelsbrunner, O'Rourke, and Seidel [EOS86] by a different method. The first correct proof of the general d-dimensional case, essentially the one presented here, is due to Edelsbrunner, Seidel, and Sharir [ESS93]. The main ingredients of the technique were previously developed by Sharir and his coauthors in several papers.

Bern, Eppstein, Plassman, and Yao [BEPY91] determined the best constant in the planar zone theorem: The zone of a line in an arrangement of n lines has at most $5.5n$ edges. They also showed that the zone of a convex k-gon has complexity $O(n + k^2)$.

The extension of the zone theorem to the zone of a k-dimensional algebraic variety in a hyperplane arrangement, as mentioned in the text, was proved by Aronov, Pellegrini, and Sharir [APS93]. They also obtained the same bound with ζ being the relative boundary of a $(k+1)$-dimensional convex set in $\mathbf{R}^d$.

The problem with the zone of a curved surface that did not exist for the zone of a hyperplane is that a face F of the zone of ζ can be split by a newly inserted hyperplane h into two subfaces F_1 and F_2, both of them lying in the zone, without $h \cap F$ being in the zone of $\zeta \cap h$, as is illustrated below:

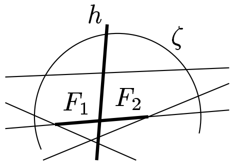

It turns out that each face F split by h in this way is adjacent to a facet in h that can be seen from ζ from both sides; such a facet is called a *popular facet* of the zone. In order to set up a suitable recurrence for the number of faces in the zone, one needs to bound the total complexity of all popular facets. This is again done by a technique similar to the proof of the zone theorem in the text. The concept of popular facet needs to be generalized to a *popular j-face*, which is a j-dimensional face F that can be seen from ζ in all the 2^{d-j} "sectors" determined by the $d-j$ hyperplanes defining F. The key observation is that if a blue popular j-face is split into two new popular j-faces by the new red hyperplane, then this can be charged to a popular $(j-1)$-face within h, as the following picture illustrates for $j=1$:

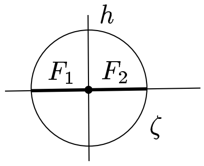

This is used to set up recurrences for the numbers of popular j-faces.

Exercises

1. (Sum of squares of cell complexities)

   (a) Let $\mathcal{C}$ be the set of all cells of an arrangement of a set H of n hyperplanes in $\mathbf{R}^d$. For $d=2,3$, prove that $\sum_{C\in\mathcal{C}} f_0(C)^2 = O(n^d)$, where $f_0(C)$ is the number of vertices of the cell C. $\boxed{3}$

   (b) Use the technique explained in this section to prove $\sum_{C\in\mathcal{C}} f_0(C)^2 = O(n^d(\log n)^{\lfloor d/2\rfloor -1})$ for every fixed $d \geq 3$ (or a similar bound with a larger constant in the exponent of $\log n$ if it helps). $\boxed{5}$

   The result in (b) is from Aronov, Matoušek, and Sharir [AMS94].

2. Define the $(\leq k)$-zone of a hyperplane g in an arrangement of hyperplanes as the collection of all faces for which some point x of their relative interior can be connected to some point $y \in g$ so that the interior of the segment xy intersects at most k hyperplanes.

   (a) By the technique of Section 6.3 (Clarkson's theorem on levels), show that the number of vertices of the $(\leq k)$-zone is $O(n^{d-1}k)$. $\boxed{3}$

   (b) Show that the bound in (a) cannot be improved in general. $\boxed{2}$

3. In this exercise we aim at bounding $K(n,n)$, the maximum total number of edges of n distinct cells in an arrangement of n lines in the plane, using the cutting lemma as in Section 4.5 (this proof is due to Clarkson, Edelsbrunner, Guibas, Sharir, and Welzl [CEG+90]). Let L be a set of n lines in general position.

   (a) Prove the bound $K(n,m) = O(n\sqrt{m}+m)$. $\boxed{3}$

   (b) Prove $K(n,n) = O(n^{4/3})$ using the cutting lemma. $\boxed{4}$

4. Consider a set H of n planes in $\mathbf{R}^3$ in general position and a sphere S (the surface of a ball).

   (a) Show that S intersects at most $O(n^2)$ cells of the arrangement of H. $\boxed{2}$

   (b) Using (a) and Exercise 1, prove that the zone of S in the arrangement of H has at most $O(n^{5/2})$ vertices. $\boxed{4}$ (This is just an upper bound; the correct order of magnitude is about n^2.)

6.5 The Cutting Lemma Revisited

Here we present the most advanced version of the random sampling technique. It combines the approach to the weak cutting lemma (Section 4.6) with ingredients from the proof of Clarkson's theorem on levels and additional ideas.

We are going to re-prove the cutting lemma 4.5.3: *For every set H of n lines in the plane and every $r > 1$ there exists a $\frac{1}{r}$-cutting for H of size $O(r^2)$, i.e., a subdivision of the plane into $O(r^2)$ generalized triangles $\Delta_1, \ldots, \Delta_t$ such that the interior of each Δ_i is intersected by at most $\frac{n}{r}$ lines of H.* The proof uses random sampling, and unlike the elementary proof in Section 4.7, it can be generalized to higher dimensions without much trouble. We first give a complete proof for the planar case and then we discuss the generalizations.

Throughout this section we assume that H is in general position. A perturbation argument mentioned in Section 4.7 can be used to derive the cutting lemma for an arbitrary H.

The first idea is as in the proof of a weaker cutting lemma by random sampling in Section 4.6: We pick a random sample S of a suitable size and triangulate its arrangement.

The subsequent calculations become simpler and more elegant if we choose S by independent Bernoulli trials. That is, instead of picking s random lines with repetitions as in Section 4.6, we fix a probability $p = \frac{s}{n}$ and we include each line $h \in H$ into S with probability p, the decisions being mutually independent (this is as in the proof of the planar case of Clarkson's theorem on levels). These two ways of random sampling (by s random draws with repetitions and by independent trials with success probability $\frac{s}{n}$) can usually be thought of as nearly the same; although the actual calculations differ significantly, their results tend to be similar.

Sampling and triangulation alone do not work. Considerations similar to those in Section 4.6 show that with probability close to 1, none of the triangles in the triangulation for the random sample S as above is intersected by more than $C\frac{n}{s}\log n$ lines of H, for a suitable constant C. Later we will see that a similar statement is true with $C\frac{n}{s}\log s$ instead of $C\frac{n}{s}\log n$. But it is *not* generally true with $C\frac{n}{s}$, for any C independent of s and n. So the most direct road to an optimal $\frac{1}{r}$-cutting, namely choosing *const* $\cdot\, r$ random lines and triangulating their arrangement, is impassable.

To see this, consider a 1-dimensional situation, where $H = \{h_1, \ldots, h_n\}$ is a set of n points in $\mathbf{R}$ (or if you prefer, look at the part of a 2-dimensional arrangement along one of the lines). For simplicity, let us set $s = \frac{n}{2}$; then $p = \frac{1}{2}$, and we can imagine that we toss a fair coin n times and we include h_i

into S if the ith toss is heads. The picture illustrates the result of 30 tosses, with black dots indicating heads:

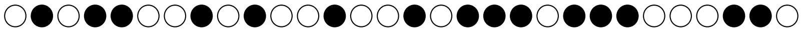

We are interested in the length of the longest consecutive run of tails (empty circles). For k is fixed, it is very likely that k consecutive tails show up in a sequence of n tosses for n sufficiently large. Indeed, if we divide the tosses into blocks of length k (suppose for simplicity that n is divisible by k),

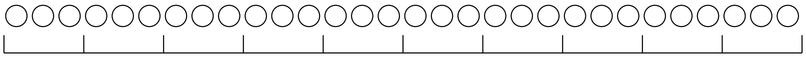

then in each block, we have probability 2^{-k} of receiving all tails. The blocks are mutually independent, and so the probability of not obtaining all tails in any of the $\frac{n}{k}$ blocks is $(1-2^{-k})^{n/k}$. For k fixed and $n \to \infty$ this goes to 0, and a more careful calculation shows that for $k = \lfloor \frac{1}{2} \log_2 n \rfloor$ we have exponentially small probability of not receiving any block of k consecutive tails (Exercise 1). So a sequence of n tosses is very likely to contain about $\log n$ consecutive tails. (Sequences produced by humans that are intended to look random usually do not have this property; they tend to be "too uniform.") Similarly, for a smaller s, if we make a circle black with probability $\frac{s}{n}$, then the longest run typically has about $\frac{n}{s} \log s$ consecutive empty circles.

Of course, in the one-dimensional situation one can define much more uniform samples, say by making every $\frac{n}{s}$th circle black. But it is not clear how one could produce such "more uniform" samples for lines in the plane or for hyperplanes in $\mathbf{R}^d$.

The strategy: a two-level decomposition. Instead of trying to select better samples we construct a $\frac{1}{r}$-cutting for H in two stages. First we take a sample S with probability $p = \frac{r}{n}$ and triangulate the arrangement, obtaining a collection $\mathcal{T}$ of triangles. (The expected number of triangles is $O(r^2)$, as we will verify later.) Typically, $\mathcal{T}$ is not yet a $\frac{1}{r}$-cutting. Let $I(\Delta)$ denote the set of lines of H intersecting the interior of a triangle $\Delta \in \mathcal{T}$ and let $n_\Delta = |I(\Delta)|$. We define the *excess* of a triangle $\Delta \in \mathcal{T}$ as $t_\Delta = n_\Delta \cdot \frac{r}{n}$.

If $t_\Delta \leq 1$, then $n_\Delta \leq \frac{n}{r}$ and Δ is a good citizen: It can be included into the final $\frac{1}{r}$-cutting as is. On the other hand, if $t_\Delta > 1$, then Δ needs further treatment: We subdivide it into a collection of finer triangles such that each of them is intersected by at most $\frac{n}{r}$ lines of H. We do it in a seemingly naive way: We consider the whole arrangement of $I(\Delta)$, temporarily ignoring Δ, and we construct a $\frac{1}{t_\Delta}$-cutting for it. Then we intersect the triangles of this $\frac{1}{t_\Delta}$-cutting with Δ, which can produce triangles but also quadrilaterals, pentagons, and hexagons. Each of these convex polygons is further subdivided into triangles, as is illustrated below:

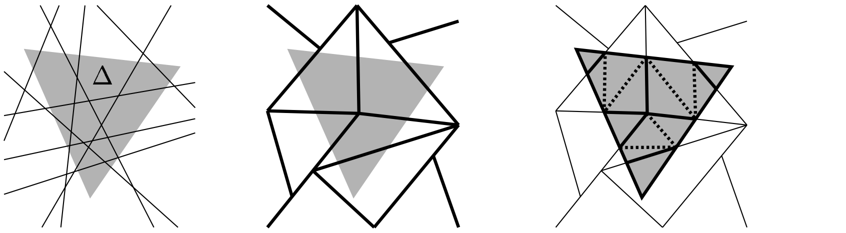

$I(\Delta)$ a $\frac{1}{t_\Delta}$-cutting restrict to Δ and triangulate

Note that each triangle in the $\frac{1}{t_\Delta}$-cutting is intersected by at most $\frac{n_\Delta}{t_\Delta} = \frac{n}{r}$ lines of $I(\Delta)$. Therefore, the triangles obtained within Δ are valid triangles of a $\frac{1}{r}$-cutting for H. The final $\frac{1}{r}$-cutting for H is constructed by subdividing each $\Delta \in \mathcal{T}$ with excess greater than 1 in the indicated manner and taking all the resulting triangles together.

How do we make the required $\frac{1}{t_\Delta}$-cuttings for the $I(\Delta)$? We do not yet have any suitable way of doing this unless we use the cutting lemma itself, which we do not want, of course. Fortunately, as a by-product of the subsequent considerations, we obtain a method for directly constructing slightly suboptimal cuttings:

6.5.1 Lemma (A suboptimal cutting lemma). *For every finite collection of lines and any $u > 1$, there exists a $\frac{1}{u}$-cutting consisting of at most $K(u\log(u{+}1))^2$ triangles, where K is a suitable constant.*

If we employ this lemma for producing the $\frac{1}{t_\Delta}$-cuttings, we can estimate the number of triangles in the resulting $\frac{1}{r}$-cutting in terms of the excesses of the triangles in $\mathcal{T}$: The total number of triangles is bounded by

$$\sum_{\Delta \in \mathcal{T}} \max\left\{1, 4K(t_\Delta \log(t_\Delta + 1))^2\right\}. \tag{6.3}$$

The key insight for the proof of the cutting lemma is that although we typically do have triangles $\Delta \in \mathcal{T}$ with excess as large as about $\log r$, they are very few. More precisely, we show that under suitable assumptions, the expected number of triangles in $\mathcal{T}$ with excess t or larger decreases exponentially as a function of t. This will take care of both estimating (6.3) by $O(r^2)$ and establishing Lemma 6.5.1.

Good and bad triangulations. Our collection $\mathcal{T}$ of triangles is obtained by triangulating the cells in the arrangement of the random sample S. Now is the time to specify how exactly the cells are triangulated, since not every triangulation works. To see this, consider a set H of n lines, each of them touching the unit circle, and let S be a random sample, again for simplicity with probability $p = \frac{1}{2}$. We have learned that such a sample is very likely to leave a gap of about $\log n$ unselected lines (as we go along the unit circle). If we maliciously triangulate the central cell in the arrangement of S by the diagonals from the vertex near such a large gap,

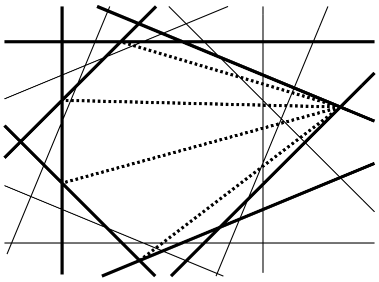

all these about $\frac{n}{2}$ triangles have excess about $\log n$; this is way too large for our purposes.

The triangulation thus cannot be quite arbitrary. For the subsequent proof, it has to satisfy simple axioms. In the planar case, it is actually technically easier not to triangulate but to construct the *vertical decomposition* of the arrangement of S. We erect vertical segments upwards and downwards from each vertex in the arrangement of S and extend them until they meet another line (or all the way to infinity):

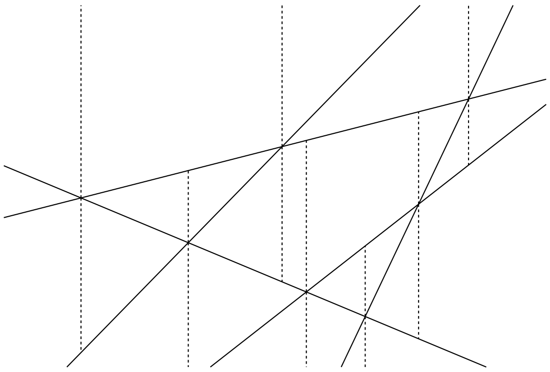

So far we have been speaking of triangles, and now we have trapezoids, but the difference is immaterial, since we can always split each trapezoid into two triangles if we wish. Let $\mathcal{T}(S)$ denote the set of (generalized) trapezoids in the vertical decomposition of S. As before, $I(\Delta)$ is the set of lines of H intersecting the interior of a trapezoid Δ.

6.5.2 Proposition (Trapezoids with large excess are rare). *Let H be a fixed set of n lines in general position, let $p = \frac{r}{n}$, where $1 \leq r \leq \frac{n}{2}$, let S be a random sample drawn from H by independent Bernoulli trials with success probability p, and let $t \geq 0$ be a real parameter. Let $\mathcal{T}(S)_{\geq t}$ denote the set of trapezoids in $\Delta \in \mathcal{T}(S)$ with excess at least t, i.e., with $|I(\Delta)| \geq t\frac{n}{r}$. Then the expected number of trapezoids in $\mathcal{T}(S)_{\geq t}$ is bounded as follows:*

$$\mathbf{E}[|\mathcal{T}(S)_{\geq t}|] \leq C \cdot 2^{-t} r^2$$

for a suitable absolute constant C.

First let us see how this result can be applied.

Proof of the suboptimal cutting lemma 6.5.1. To obtain a $\frac{1}{u}$-cutting for H, we set $r = Au\log(u+1)$ for a sufficiently large constant A and choose a sample S as in Proposition 6.5.2.

By that proposition with $t = 0$, we have $\mathbf{E}[|\mathcal{T}(S)|] \leq Br^2$ for a suitable constant B. By the same proposition with $t = A\log(u+1)$, we have $\mathbf{E}[|\mathcal{T}(S)_{\geq t}|] \leq \frac{1}{3}$ if A is sufficiently large. By linearity of expectation, we obtain

$$\mathbf{E}\left[\frac{1}{3Br^2}\,|\mathcal{T}(S)| + |\mathcal{T}(S)_{\geq t}|\right] \leq \frac{2}{3}.$$

So there exists a sample S with both $|\mathcal{T}(S)| \leq 2Br^2$ and $|\mathcal{T}(S)_{\geq t}| = 0$. This means that we have a $\frac{1}{u}$-cutting into $O(r^2) = O((u\log(u+1))^2)$ trapezoids. □

For an alternative proof of Lemma 6.5.1 see Exercise 10.3.4.

Proof of the cutting lemma (Lemma 4.5.3). Most of the proof has already been described. To produce a $\frac{1}{r}$-cutting, we pick a random sample S with probability $p = \frac{r}{n}$, we let $\mathcal{T} = \mathcal{T}(S)$ be its vertical decomposition, and we refine each trapezoid $\Delta \in \mathcal{T}$ with excess $t_\Delta > 1$ using an auxiliary $\frac{1}{t_\Delta}$-cutting. The size of the resulting $\frac{1}{r}$-cutting is bounded by (6.3). So it suffices to estimate the expected value of that expression using Proposition 6.5.2:

$$\begin{aligned}
&\mathbf{E}\Bigg[\sum_{\Delta\in\mathcal{T}(S)} \max\left\{1, 4K(t_\Delta\log(t_\Delta+1))^2\right\}\Bigg]\\
&\leq \mathbf{E}\Bigg[\sum_{\Delta\in\mathcal{T}(S)} \max\left\{1, 4Kt_\Delta^4\right\}\Bigg] \qquad\qquad \text{(as } \log(t_\Delta+1) \leq t_\Delta\text{)}\\
&\leq \mathbf{E}\Bigg[|\mathcal{T}(S)| + \sum_{i=0}^{\infty} \sum_{\substack{\Delta\in\mathcal{T}(S)\\ 2^i \leq t_\Delta < 2^{i+1}}} 4Kt_\Delta^4\Bigg]\\
&\leq \mathbf{E}[|\mathcal{T}(S)|] + \sum_{i=0}^{\infty} \mathbf{E}\left[|\mathcal{T}(S)|_{\geq 2^i}\right] \cdot O(2^{4(i+1)})\\
&\leq O(r^2) + \sum_{i=0}^{\infty} C\cdot 2^{-2^i} r^2 \cdot O(2^{4(i+1)})\\
&= O(r^2).
\end{aligned}$$

The cutting lemma is proved. □

Note that it was not important that the suboptimal cutting lemma is near-optimal: *Any bound subexponential in* u for the size of a $\frac{1}{u}$-cutting would do. In particular, for any fixed $c \geq 1$, the expected cth-degree average of the excess is only a constant.

For the proof of Proposition 6.5.2, we need several definitions and some simple properties of the vertical decomposition. Let H be a fixed set of lines in general position, and let $Reg = \bigcup_{S\subseteq H} \mathcal{T}(S)$ be the set of all trapezoids that can ever appear in the vertical decomposition for some $S \subseteq H$ (including $S = \emptyset$). For a trapezoid $\Delta \in Reg$, let $D(\Delta)$ be the set of the lines of H incident to at least one vertex of Δ. By the general-position assumption, we have $|D(\Delta)| \leq 4$ for all Δ. The various possible cases, up to symmetry, are drawn below; the picture shows the lines of $D(\Delta)$ with Δ marked in gray:

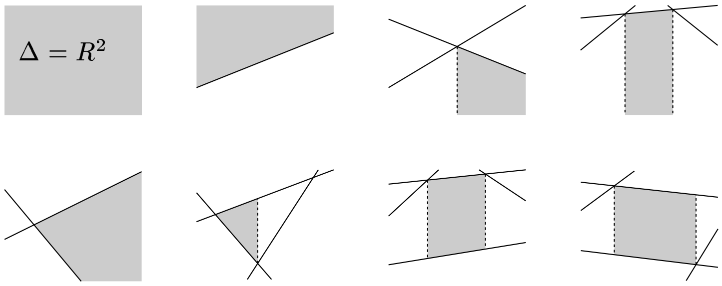

The set $D(\Delta)$ is called the *defining set* of Δ. Note that the same defining set can belong to several trapezoids.

Now we list the properties required for the proof; some of them are obvious or have already been noted.

(C0) We have $|D(\Delta)| \leq 4$ for all $\Delta \in \mathit{Reg}$. Moreover, any set $S_0 \subseteq H$ is the defining set for at most a constant number of $\Delta \in \mathit{Reg}$ (certainly no more than the maximum of $|\mathcal{T}(S_0)|$ for $|S_0| \leq 4$).

(C1) For any $\Delta \in \mathcal{T}(S)$, we have $D(\Delta) \subseteq S$ (the defining set must be present) and $S \cap I(\Delta) = \emptyset$ (no intersecting line may be present).

(C2) For any $\Delta \in \mathit{Reg}$ and any $S \subseteq H$ such that $D(\Delta) \subseteq S$ and $I(\Delta) \cap S = \emptyset$, we have $\Delta \in \mathcal{T}(S)$.

(C3) For every $S \subseteq H$, we have $|\mathcal{T}(S)| = O(|S|^2 + 1)$. To see this, think of adding the vertical segments to the arrangement of S one by one: Each of them splits an existing region in two.

The most interesting condition is (C2), which says that the vertical decomposition is defined "locally." It implies, in particular, that Δ is one of the trapezoids in the vertical decomposition of its defining set. More generally, it says that $\Delta \in \mathit{Reg}$ is present in $\mathcal{T}(S)$ whenever it is not excluded for simple local reasons (which can be checked by looking only at Δ). Checking (C2) in our situation is easy, and we leave it to the reader. Also note that it is (C2) that is generally violated for the mischievous triangulation considered earlier.

Proof of Proposition 6.5.2. First we prove that if $S \subseteq H$ is a random sample drawn with probability $p = \frac{r}{n}$, $0 < r < n$, then

$$\mathbf{E}[|\mathcal{T}(S)|] = O(r^2 + 1). \tag{6.4}$$

This takes care of the case $t \leq 1$ in the proposition. By (C3), we have $|\mathcal{T}(S)| = O(|S|^2+1)$ for every fixed S, and so it suffices to show that $\mathbf{E}\left[|S|^2\right] = O(r^2+1)$. Now, $|S|$ is the sum of independent random variables, each of them attaining value 1 with probability p and value 0 with probability $1-p$, and it is easy to check that $\mathbf{E}\left[|S|^2\right] \leq r^2 + r$ (Exercise 2(a)).

Next, we assume $t \geq 1$. Let $S \subseteq H$ be a random sample drawn with probability p. We observe that the conditions (C1) and (C2) allow us to express the probability $p(\Delta)$ that a certain trapezoid $\Delta \in Reg$ appears in the vertical decomposition $\mathcal{T}(S)$: Since Δ appears if and only if all lines of $D(\Delta)$ are selected into S and none of $I(\Delta)$ is selected, we have

$$p(\Delta) = p^{|D(\Delta)|}(1-p)^{|I(\Delta)|}.$$

(An analogous formula appeared in the proof of the planar Clarkson's theorem on levels, and one can say that the technique of that proof is developed one step further in the present proof.) If we write $Reg_{\geq t} = \{\Delta \in Reg\colon |I(\Delta)| \geq t\frac{n}{r}\}$ for the set of all potential trapezoids with excess at least t, the expected number of trapezoids in $\mathcal{T}(S)_{\geq t}$ can be written as

$$\mathbf{E}[|\mathcal{T}(S)_{\geq t}|] = \sum_{\Delta \in Reg_{\geq t}} p(\Delta). \tag{6.5}$$

It seems difficult to estimate this sum directly; the trick is to compare it with a similar sum obtained for the expected number of trapezoids for another sample.

We define another probability $\tilde{p} = \frac{p}{t}$, and we let $\tilde{S}$ be a sample drawn from H by Bernoulli trials with success probability $\tilde{p}$. On the one hand, we have $\mathbf{E}\left[|\mathcal{T}(\tilde{S})|\right] = O(r^2/t^2 + 1)$ by (6.4). On the other hand, setting $\tilde{p}(\Delta) = \tilde{p}^{|D(\Delta)|}(1-\tilde{p})^{|I(\Delta)|}$ we have, in analogy to (6.5),

$$\begin{aligned} \mathbf{E}\left[|\mathcal{T}(\tilde{S})|\right] &= \sum_{\Delta \in Reg} \tilde{p}(\Delta) \geq \sum_{\Delta \in Reg_{\geq t}} \tilde{p}(\Delta) \\ &= \sum_{\Delta \in Reg_{\geq t}} p(\Delta) \cdot \frac{\tilde{p}(\Delta)}{p(\Delta)} \geq \mathbf{E}[|\mathcal{T}(S)_{\geq t}|] \cdot R, \end{aligned} \tag{6.6}$$

where

$$R = \min\left\{\frac{\tilde{p}(\Delta)}{p(\Delta)}\colon \Delta \in Reg_{\geq t}\right\}.$$

Now R can be bounded from below. For every $\Delta \in Reg_{\geq t}$, we have $|I(\Delta)| \geq t\frac{r}{n}$ and $|D(\Delta)| \leq 4$, and so

$$\frac{\tilde{p}(\Delta)}{p(\Delta)} = \left(\frac{\tilde{p}}{p}\right)^{|D(\Delta)|}\left(\frac{1-\tilde{p}}{1-p}\right)^{|I(\Delta)|} \geq t^{-4}\left(\frac{1-\tilde{p}}{1-p}\right)^{tn/r}.$$

We use $1-p \leq e^{-p}$ (this holds for all real p) and $1-\tilde{p} \geq e^{-2\tilde{p}}$ (this is true for all $\tilde{p} \in [0, \frac{1}{2}]$, and we have $\tilde{p} \leq p \leq \frac{1}{2}$). Therefore $R \geq t^{-4}e^{t-2}$. Substituting into (6.6), we finally derive

$$\mathbf{E}[|\mathcal{T}(S)_{\geq t}|] \leq \frac{1}{R} \cdot \mathbf{E}\left[|\mathcal{T}(\tilde{S})|\right] \leq t^4 e^{-(t-2)} \cdot O\left(\frac{r^2}{t^2} + 1\right) \leq C \cdot 2^{-t} r^2$$

for a sufficiently large constant C (the proposition assumes $r \geq 1$). Proposition 6.5.2 is proved. $\square$

The following can be proved by the same technique:

6.5.3 Theorem (Cutting lemma for arbitrary dimension). *Let $d \geq 1$ be a fixed integer, let H be a set of n hyperplanes in $\mathbf{R}^d$, and let r be a parameter, $1 < r \leq n$. Then there exists a $\frac{1}{r}$-cutting for H of size $O(r^d)$; that is, a subdivision of $\mathbf{R}^d$ into $O(r^d)$ generalized simplices such that the interior of each simplex is intersected by at most $\frac{n}{r}$ hyperplanes of H.*

The only new part of the proof is the construction of a suitable triangulation scheme that plays the role of $\mathcal{T}(S)$. A vertical decomposition does not work. More precisely, it is not known whether the vertical decomposition of an arrangement of n hyperplanes in $\mathbf{R}^d$ always has at most $O(n^d)$ cells (prisms); this would be needed as the analogue of condition (C3). Instead one can use the *bottom-vertex triangulation*, which we define next.

First we specify the bottom-vertex triangulation of a k-dimensional convex polytope $P \subset \mathbf{R}^d$, $1 \leq k \leq d$, by induction on k. For $k = 1$, P is a line segment, and the triangulation consists of P itself. For $k > 1$, we let v be the vertex of P with the smallest last coordinate (the "bottom vertex"); ties can be broken by lexicographic ordering of the coordinate vectors. We triangulate all proper faces of P inductively, and we add the simplices obtained by erecting the cone with apex v over all simplices in the triangulations of the faces not containing v.

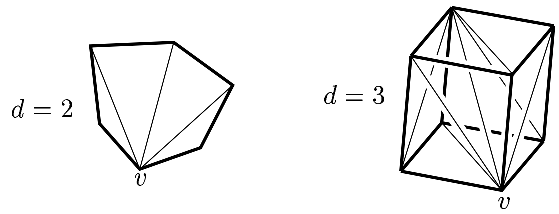

It is not difficult to check that this yields a triangulation of P (even a simplicial complex, although this is not needed in the present proof), and that if P is a simple polytope, then the total number of simplices in this triangulation is at most proportional to the number of vertices of P (with the constant of proportionality depending on d); see Exercise 4.

All the bounded cells of the arrangement of S are triangulated in this way. Some care is needed for the unbounded cells, and several ways are available. One of the simplest is to intersect the arrangement with a sufficiently large box containing all the vertices and construct the $\frac{1}{r}$-cutting only inside that box. This is sufficient for most applications of $\frac{1}{r}$-cuttings. Alternatively (and almost equivalently), we can consider the whole arrangement in the projective d-space instead of $\mathbf{R}^d$. We omit a detailed discussion of this aspect.

In this way we obtain a triangulation $\mathcal{T}(S)$ for every subset S of the given set of hyperplanes. The analogue of (C3) is $|\mathcal{T}(S)| = O(|S|^d + 1)$, which follows (assuming H in general position) because the number of simplices in each cell is proportional to the number of its vertices, and the total number of vertices is $O(|S|^d)$.

The set $I(\Delta)$ are all hyperplanes intersecting the interior of a simplex Δ, and $D(\Delta)$ consists of all the hyperplanes incident to at least one vertex of Δ. We again need to assume that our hyperplanes are in general position.

Then, obviously, $|D(\Delta)| \leq d(d+1)$, and a more careful argument shows that $|D(\Delta)| \leq \frac{d(d+3)}{2}$. The important thing is that an analogue of (C0) holds, namely, that both $|D(\Delta)|$ and the number of Δ with a given $D(\Delta)$ are bounded by constants.

The condition (C1) holds trivially. The "locality" condition (C2) does need some work, although it is not too difficult, and we refer to Chazelle and Friedman [CF90] for a detailed argument.

With (C0)–(C3) in place, the whole proof proceeds exactly as in the planar case. To get the analogue of (6.4), namely $\mathbf{E}[|\mathcal{T}(S)|] = O(r^d+1)$, we need the fact that $\mathbf{E}\left[|S|^d\right] = O(r^d)$ (this is what we avoided in the proof of the higher-dimensional Clarkson's theorem on levels by passing to another way of sampling); see Exercise 2(b) or 3.

Further generalizations. An analogue of Proposition 6.5.2 can be derived from conditions (C0)–(C3) in a general abstract framework. It provides optimal $\frac{1}{r}$-cuttings not only for arrangements of hyperplanes but also in other situations, whenever one can define a suitable decomposition scheme satisfying (C0)–(C3) and bound the maximum number of cells in the decomposition (the latter is a challenging open problem for arrangements of bounded-degree algebraic surfaces). The significance of Proposition 6.5.2 reaches beyond the construction of cuttings; its variations have been used extensively, mainly in the analysis of geometric algorithms. We are going to encounter a combinatorial application in Chapter 11.

Bibliography and remarks. The proof of the cutting lemma as in this section (with a different way of sampling) is due to Chazelle and Friedman [CF90]. Analogues of Proposition 6.5.2, or more precisely the consequence stating that the expectation of the cth degree average of the excess is bounded by a constant, were first proved and applied by Clarkson [Cla88a] (see Clarkson and Shor [CS89] for the journal version). Since then, they became one of the indispensable tools in the analysis of randomized geometric algorithms, as is illustrated by the book by Mulmuley [Mul93a], for example, as well as by many newer papers.

The bottom-vertex triangulation (also called the *canonical triangulation* in some papers) was defined in Clarkson [Cla88b].

Proposition 6.5.2 can be formulated and proved in an abstract framework, where H and Reg are some finite sets and $\mathcal{T}\colon 2^H \to 2^{Reg}$, $I\colon Reg \to 2^H$, and $D\colon Reg \to 2^H$ are mappings satisfying (C0) (with some constants), (C1), (C2), and an analogue of (C3) that bounds the expected size of $\mathcal{T}(S)$ for a random $S \subseteq H$ by a suitable function of r, typically by $O(r^k)$ for some real constant $k \geq 1$. The conclusion is $\mathbf{E}[|\mathcal{T}(S)_{\geq t}|] = O(2^{-t}r^k)$. Very similar abstract frameworks are discussed in Mulmuley [Mul93a] and in De Berg, Van Kreveld, Overmars, and Schwarzkopf [dBvKOS97].

The axiom (C2) can be weakened to the following:

(C2′) If $\Delta \in \mathcal{T}(S)$ and $S' \subseteq S$ satisfies $D(\Delta) \subseteq S'$, then $\Delta \in \mathcal{T}(S')$. That is, Δ cannot be destroyed by deleting elements of S unless we delete an element of $D(\Delta)$.

A typical situation where (C2′) holds although (C2) fails is that in which H is a set of lines in the plane and $\mathcal{T}(S)$ are the trapezoids in the vertical decomposition of the cell in the arrangement of S that contains some fixed point, say 0. Then Δ can be made to disappear by adding a line to S even if that line does not intersect

Δ, as is illustrated below:

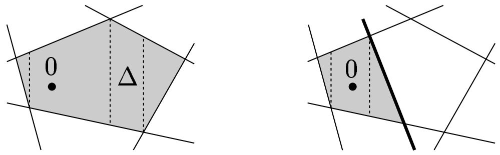

This weaker axiom was first used instead of (C2) by Chazelle, Edelsbrunner, Guibas, Sharir, and Snoeyink [CEG+93]. For a proof of a counterpart of Proposition 6.5.2 under (C2′) see Agarwal, Matoušek, and Schwarzkopf [AMS98].

Yet another proof of the cutting lemma in arbitrary dimension was invented by Chazelle [Cha93a]. An outline of the argument can also be found in Chazelle's book [Cha00c] or in the chapter by Matoušek in [SU00].

Both the proofs of the higher-dimensional cutting lemma depend crucially on the fact that the arrangement of n hyperplanes in $\mathbf{R}^d$, d fixed, can be triangulated using $O(n^d)$ simplices. As was explained in Section 6.2, the arrangement of n bounded-degree algebraic surfaces in $\mathbf{R}^d$ has $O(n^d)$ faces in total, but the faces can be arbitrarily complicated. A challenging open problem is whether each face can be further decomposed into "simple" pieces (each of them defined by a constant-bounded number of bounded-degree algebraic inequalities) such that the total number of pieces for the whole arrangement is $O(n^d)$ or not much larger. This is easy for $d = 2$ (the vertical decomposition will do), but dimension 3 is already quite challenging. Chazelle, Edelsbrunner, Guibas, and Sharir [CEGS89] found a general argument which provides an $O(n^{2d-2})$ bound in dimension d using a suitable vertical decomposition. By proving a near-optimal bound in the 3-dimensional case and using it as a basis of the induction, they obtained the bound of $O(n^{2d-3}\beta(n))$, where β is a very slowly growing function (much smaller than $\log^* n$). Recently Koltun [Kol01] established a near-tight bound in the 4-dimensional situation, which pushed the general bound to $O(n^{2d-4+\varepsilon})$ for every fixed $d \geq 4$.

This decomposition problem is the main obstacle to proving an optimal or near-optimal cutting lemma for arrangements of algebraic surfaces. For some special cases, say for an arrangement of spheres in $\mathbf{R}^d$, optimal decompositions are known and an optimal cutting lemma can be obtained. In general, if one can prove a bound of $O(n^\alpha)$ for the number of pieces in the decomposition, then the techniques of Chapter 10 yield $\frac{1}{r}$-cuttings of size $O(r^\alpha \log^\alpha r)$, and if, moreover, the locality condition (C2) can be guaranteed, then the method of the present section leads to $\frac{1}{r}$-cuttings of size $O(r^\alpha)$.

Exercises

1. Estimate the largest $k = k(n)$ such that in a row of n tosses of a fair coin we obtain k consecutive tails with probability at least $\frac{1}{2}$. In particular, using the trick with blocks in the text, check that for $k = \lfloor \frac{1}{2} \log_2 n \rfloor$, the probability of not getting all tails in any of the blocks is exponentially small (as a function of n). $\boxed{2}$

2. Let $X = X_1 + X_2 + \cdots + X_n$, where the X_i are independent random variables, each attaining the value 1 with probability p and the value 0 with probability $1 - p$.

   (a) Calculate $\mathbf{E}\left[X^2\right]$. $\boxed{2}$

   (b) Prove that for every integer $d \geq 1$ there exists c_d such that $\mathbf{E}\left[X^d\right] \leq (np + c_d)^d$. (You can use a Chernoff-type inequality, or prove by induction that

$\mathbf{E}\left[(X+a)^d\right] \leq (np+d+a)^d$ for all nonnegative integers n, d, and a.) $\boxed{3}$

(c) Use (b) to prove that $\mathbf{E}[X^\alpha] \leq (np+c_\alpha)^\alpha$ also holds for nonintegral $\alpha \geq 1$. $\boxed{3}$

3. Let $X = X_1 + X_2 + \cdots + X_n$ be as in the previous exercise. Show that $\mathbf{E}\left[\binom{X}{d}\right] = p^d\binom{n}{d}$ (where $d \geq 0$ is an integer) and conclude that $\mathbf{E}\left[X^d\right] \leq c_d(np)^d$ for $np \geq d$ and a suitable $c_d > 0$. $\boxed{3}$

4. Let P be a d-dimensional simple convex polytope. Prove that the bottom-vertex triangulation of P has at most $C_d f_0(P)$ simplices, where C_d depends only on d and $f_0(P)$ denotes the number of vertices of P. $\boxed{4}$