hence $u_j \in \operatorname{cone}(A_j)$. Above we have derived $\sum_{j=1}^{r} \varphi_j(u_j) = 0$, and so by (8.1) we get $u_1 = u_2 = \cdots = u_r$. Hence the common value of all the u_j belongs to $\bigcap_{j=1}^{r} \operatorname{cone}(A_j)$.

It remains to check that $u_j \neq 0$. Since we assume $0 \notin \operatorname{conv}(A)$, the only nonnegative linear combination of points of A equal to 0 is the trivial one, with all coefficients 0. On the other hand, since not all the α_i are 0, at least one u_j is expressed as a nontrivial linear combination of points of A. This proves Proposition 8.3.2 and Tverberg's theorem as well. □

The colored Tverberg theorem. If we have 9 points in the plane, 3 of them red, 3 blue, and 3 white, it turns out that we can always partition them into 3 triples in such a way that each triple has one red, one blue, and one white point, and the 3 triangles determined by the triples have a nonempty intersection.

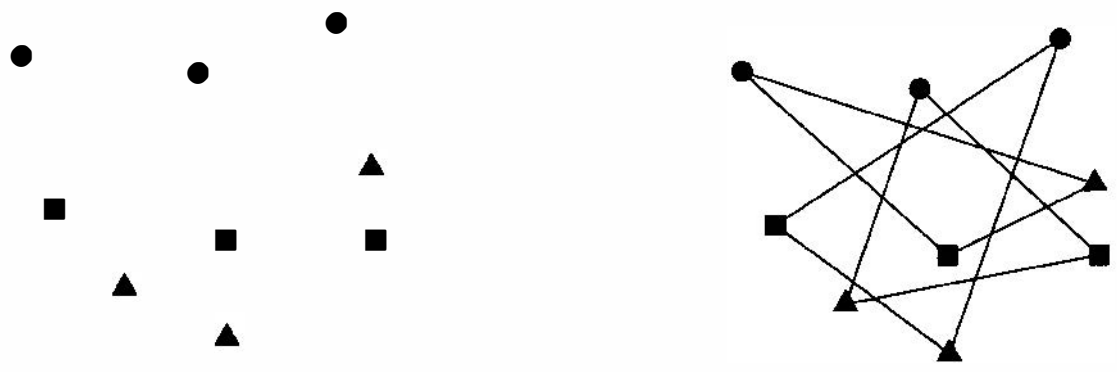

The colored Tverberg theorem is a generalization of this statement for arbitrary d and r. We will need it in Section 9.2, for a result about many simplices with a common point. In that application, the colored version is essential (and Tverberg's theorem alone is not sufficient).

8.3.3 Theorem (Colored Tverberg theorem). *For any integers $r, d \geq 2$ there exists an integer t such that given any $t(d+1)$-point set $Y \subset \mathbf{R}^d$ partitioned into $d+1$ color classes $Y_1, \ldots, Y_{d+1}$ with t points each, there exist r pairwise disjoint sets $A_1, \ldots, A_r$ such that each A_i contains exactly one point of each Y_j, $j = 1, 2, \ldots, d+1$ (that is, the A_i are rainbow), and $\bigcap_{i=1}^{r} \operatorname{conv}(A_i) \neq \emptyset$.*

Let $T_{\mathrm{col}}(d, r)$ denote the smallest t for which the conclusion of the theorem holds. It is known that $T_{\mathrm{col}}(2, r) = r$ for all r. It is possible that $T_{\mathrm{col}}(d, r) = r$ for all d and r, but only weaker bounds have been proved. The strongest known result guarantees that $T_{\mathrm{col}}(d, r) \leq 2r-1$ whenever r is a prime power.

Recall that in Tverberg's theorem, if we need only the existence of $T(d, r)$, rather than the precise value, several simple arguments are available. In contrast, for the colored version, even if we want only the existence of $T_{\mathrm{col}}(d, r)$, there is essentially only one type of proof, which is not easy and which uses topological methods. Since such methods are not considered in this book, we have to omit a proof of the colored Tverberg theorem.

Bibliography and remarks. Tverberg's theorem was conjectured by Birch and proved by Tverberg (really!) [Tve66]. His original proof is

$a \in \mathbf{R}^d$ is contained in more than dn^{d-1} hyperplanes of $\mathcal{H}$. Consequently, at most $O(n^d)$ X-simplices have a on their boundary.

Proof. For each d-tuple S whose hyperplane contains a, we choose an inclusion-minimal set $K(S) \subseteq S$ whose affine hull contains a. We claim that if $|K(S_1)| = |K(S_2)| = k$, then either $K(S_1) = K(S_2)$ or $K(S_1)$ and $K(S_2)$ share at most $k-2$ points.

Indeed, if $K(S_1) = \{x_1, \ldots, x_{k-1}, x_k\}$ and $K(S_2) = \{x_1, \ldots, x_{k-1}, y_k\}$, $x_k \neq y_k$, then the affine hulls of $K(S_1)$ and $K(S_2)$ are distinct, for otherwise, we would have $k+1$ points in a common $(k-1)$-flat, contradicting the general position of X. But then the affine hulls intersect in the $(k-2)$-flat generated by $x_1, \ldots, x_{k-1}$ and containing a, and $K(S_1)$ and $K(S_2)$ are not inclusion-minimal.

Therefore, the first $k-1$ points of $K(S)$ determine the last one uniquely, and the number of distinct sets of the form $K(S)$ of cardinality k is at most n^{k-1}. The number of hyperplanes determined by X and containing a given k-point set $K \subset X$ is at most n^{d-k}, and the lemma follows by summing over k. □

Bibliography and remarks. The planar version of the first selection lemma, with the best possible constant $\frac{2}{9}$, was proved by Boros and Füredi [BF84]. A generalization to an arbitrary dimension, with the first of the two proofs given above, was found by Bárány [Bár82]. The idea of the proof of Lemma 9.1.2 was communicated to me by János Pach.

Boros and Füredi [BF84] actually showed that any centerpoint of X works; that is, it is contained in at least $\frac{2}{9}\binom{n}{3}$ X-triangles. Wagner and Welzl (private communication) observed that a centerpoint works in every fixed dimension, being common to at least $c_d\binom{n}{d+1}$ X-simplices. This follows from known results on the face numbers of convex polytopes using the Gale transform, and it provides yet another proof of the first selection lemma, yielding a slightly better value of the constant c_d than that provided by Bárány's proof. Moreover, for a centrally symmetric point set X this method implies that the origin is contained in the largest possible number of X-simplices.

As for lower bounds, it is known that no n-point $X \subset \mathbf{R}^d$ in general position has a point common to more than $\frac{1}{2^d}\binom{n}{d+1}$ X-simplices [Bár82]. It seems that suitable sets might provide stronger lower bounds, but no results in this direction are known.

9.2 The Second Selection Lemma

In this section we continue using the term X-simplex in the sense of Section 9.1; that is, an X-simplex is the convex hull of a $(d+1)$-point subset

of X. In that section we saw that if X is a set in $\mathbf{R}^d$ and we consider *all* the X-simplices, then at least a fixed fraction of them have a point in common.

What if we do not have all, but many X-simplices, some α-fraction of all? It turns out that still many of them must have a point in common, as stated in the second selection lemma below.

9.2.1 Theorem (Second selection lemma). *Let X be an n-point set in $\mathbf{R}^d$ and let $\mathcal{F}$ be a family of $\alpha\binom{n}{d+1}$ X-simplices, where $\alpha \in (0,1]$ is a parameter. Then there exists a point contained in at least*

$$c\alpha^{s_d}\binom{n}{d+1}$$

X-simplices of $\mathcal{F}$, where $c = c(d) > 0$ and s_d are constants.

This result is already interesting for α fixed. But for the application that motivated the discovery of the second selection lemma, namely, trying to bound the number of k-sets (see Chapter 11), the dependence of the bound on α is important, and it would be nice to determine the best possible values of the exponent s_d.

For $d = 1$ it is not too difficult to obtain an asymptotically sharp bound (see Exercise 1). For $d = 2$ the best known bound (probably still not sharp) is as follows: If $|\mathcal{F}| = n^{3-\nu}$, then there is a point contained in at least $\Omega(n^{3-3\nu}/\log^5 n)$ X-triangles of $\mathcal{F}$. In the parameterization as in Theorem 9.2.1, this means that s_2 can be taken arbitrarily close to 3, provided that α is sufficiently small, say $\alpha \le n^{-\delta}$ for some $\delta > 0$. For higher dimensions, the best known proof gives $s_d \approx (4d+1)^{d+1}$.

Hypergraphs. It is convenient to formulate some of the subsequent considerations in the language of hypergraphs. Hypergraphs are a generalization of graphs where edges can have more than 2 points (from another point of view, a hypergraph is synonymous with a set system). A *hypergraph* is a pair $H = (V, E)$, where V is the vertex set and $E \subseteq 2^V$ is a system of subsets of V, the edge set. A *k-uniform hypergraph* has all edges of size k (so a graph is a 2-uniform hypergraph). A *k-partite hypergraph* is one where the vertex set can be partitioned into k subsets $V_1, V_2, \ldots, V_k$, the *classes*, so that each edge contains at most one point from each V_i. The notions of *subhypergraph* and *isomorphism* are defined analogously to these for graphs. A subhypergraph is obtained by deleting some vertices and some edges (all edges containing the deleted vertices, but possibly more). An isomorphism is a bijection of the vertex sets that maps edges to edges in both directions (a renaming of the vertices).

Proof of the second selection lemma. The proof is somewhat similar to the second proof of the first selection lemma (Theorem 9.1.1). We again use the fractional Helly theorem. We need to show that many $(d+1)$-tuples of X-simplices of $\mathcal{F}$ are good (have nonempty intersections).

We can view $\mathcal{F}$ as a $(d+1)$-uniform hypergraph. That is, we regard X as the vertex set and each X-simplex corresponds to an edge, i.e., a subset of X of size $d+1$. This hypergraph captures the "combinatorial type" of the family $\mathcal{F}$, and a specific placement of the points of X in $\mathbf{R}^d$ then gives a concrete "geometric realization" of $\mathcal{F}$.

First, let us concentrate on the simpler task of exhibiting at least one good $(d+1)$-tuple; even this seems quite nontrivial. Why cannot we proceed as in the second proof of the first selection lemma? Let us give a concrete example with $d = 2$. Following that proof, we would consider 9 points in $\mathbf{R}^2$, and Tverberg's theorem would provide a partition into triples with intersecting convex hulls:

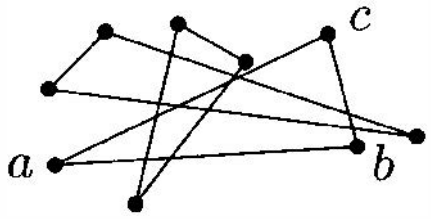

But it can easily happen that one of these triples, say $\{a, b, c\}$, is not an edge of our hypergraph. Tverberg's theorem gives us no additional information on which triples appear in the partition, and so this argument would guarantee a good triple only if *all* the triples on the considered 9 points were contained in $\mathcal{F}$. Unfortunately, a 3-uniform hypergraph on n vertices can contain more than half of all possible $\binom{n}{3}$ triples without containing all triples on some 9 points (even on 4 points). This is a "higher-dimensional" version of the fact that the complete bipartite graph on $\frac{n}{2} + \frac{n}{2}$ vertices has about $\frac{1}{4}n^2$ edges without containing a triangle.

Hypergraphs with many edges need not contain complete hypergraphs, but they have to contain complete multipartite hypergraphs. For example, a graph on n vertices with significantly more than $n^{3/2}$ edges contains $K_{2,2}$, the complete bipartite graph on $2 + 2$ vertices (see Section 4.5). Concerning hypergraphs, let $K^{d+1}(t)$ denote the complete $(d+1)$-partite $(d+1)$-uniform hypergraph with t vertices in each of its $d+1$ vertex classes. The illustration shows a $K^3(4)$; only three edges are drawn as a sample, although of course, all triples connecting vertices at different levels are present.

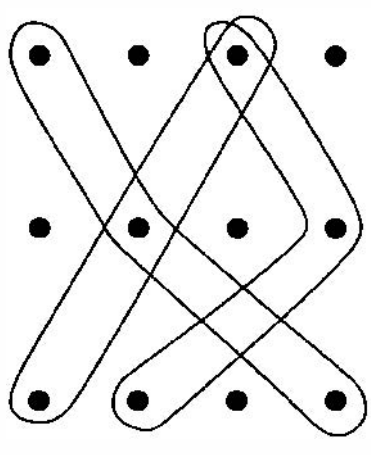

If t is a constant and we have a $(d+1)$-uniform hypergraph on n vertices with sufficiently many edges, then it has to contain a copy of $K^{d+1}(t)$ as a subhypergraph. We do not formulate this result precisely, since we will need a stronger one later.

In geometric language, given a family $\mathcal{F}$ of sufficiently many X-simplices, we can color some t points of X red, some other t points blue,..., t points by color $(d+1)$ in such a way that all the rainbow X-simplices on the $(d+1)t$ colored points are present in $\mathcal{F}$. And in such a situation, if t is a sufficiently large constant, the colored Tverberg theorem (Theorem 8.3.3) with $r = d+1$ claims that we can find a $(d+1)$-tuple of vertex-disjoint rainbow X-simplices whose convex hulls intersect, and so there is a good $(d+1)$-tuple! In fact, these are the considerations that led to the formulation of the colored Tverberg theorem.

For the fractional Helly theorem, we need not only one but many good $(d+1)$-tuples. We use an appropriate stronger hypergraph result, saying that if a hypergraph has enough edges, then it contains many copies of $K^{d+1}(t)$:

9.2.2 Theorem (The Erdős–Simonovits theorem). *Let d and t be positive integers. Let $\mathcal{H}$ be a $(d+1)$-uniform hypergraph on n vertices and with $\alpha\binom{n}{d+1}$ edges, where $\alpha \geq Cn^{-1/t^d}$ for a certain sufficiently large constant C. Then $\mathcal{H}$ contains at least*

$$c\alpha^{t^{d+1}} n^{(d+1)t}$$

copies of $K^{d+1}(t)$, where $c = c(d,t) > 0$ is a constant.

For completeness, a proof is given at the end of this section.

Note that in particular, the theorem implies that a $(d+1)$-uniform hypergraph having at least a constant fraction of all possible edges contains at least a constant fraction of all possible copies of $K^{d+1}(t)$.

We can now finish the proof of the second selection lemma by double counting. The given family $\mathcal{F}$, viewed as a $(d+1)$-uniform hypergraph, has $\alpha\binom{n}{d+1}$ edges, and thus it contains at least $c\alpha^{t^{d+1}} n^{(d+1)t}$ copies of $K^{d+1}(t)$ by Theorem 9.2.2. As was explained above, each such copy contributes at least one good $(d+1)$-tuple of vertex-disjoint X-simplices of $\mathcal{F}$. On the other hand, $d+1$ vertex-disjoint X-simplices have together $(d+1)^2$ vertices, and hence their vertex set can be extended to a vertex set of some $K^{d+1}(t)$ (which has $t(d+1)$ vertices) in at most $n^{t(d+1)-(d+1)^2} = n^{(t-d-1)(d+1)}$ ways. This is the maximum number of copies of $K^{d+1}(t)$ that can give rise to the same good $(d+1)$-tuple. Hence there are at least $c\alpha^{t^{d+1}} n^{(d+1)^2}$ good $(d+1)$-tuples of X-simplices of $\mathcal{F}$. By the fractional Helly theorem, at least $c'\alpha^{t^{d+1}} n^{d+1}$ X-simplices of $\mathcal{F}$ share a common point, with $c' = c'(d) > 0$. This proves the second selection lemma, with the exponent $s_d \leq (4d+1)^{d+1}$. □

Proof of the Erdős–Simonovits theorem (Theorem 9.2.2). By induction on k, we are going to show that a k-uniform hypergraph on n vertices and with m edges contains at least $f_k(n,m)$ copies of $K^k(t)$, where

$$f_k(n,m) = c_k n^{tk} \left(\frac{m}{n^k}\right)^{t^k} - C_k n^{t(k-1)},$$

with $c_k > 0$ and C_k suitable constants depending on k and also on t (t is not shown in the notation, since it remains fixed). This claim with $k = d+1$ implies the Erdős–Simonovits theorem.

For $k = 1$, the claim holds.

So let $k > 1$ and let $\mathcal{H}$ be k-uniform with vertex set V, $|V| = n$, and edge set E, $|E| = m$. For a vertex $v \in V$, define a $(k-1)$-uniform hypergraph $\mathcal{H}_v$ on V, whose edges are all edges of $\mathcal{H}$ that contain v, but with v deleted; that is, $\mathcal{H}_v = (V, \{e \setminus \{v\}: e \in E, v \in e\})$. Further, let $\mathcal{H}'$ be the $(k-1)$-uniform hypergraph whose edge set is the union of the edge sets of all the $\mathcal{H}_v$.

Let $\mathcal{K}$ denote the set of all copies of the complete $(k-1)$-partite hypergraph $K^{k-1}(t)$ in $\mathcal{H}'$. The key notion in the proof is that of an *extending vertex* for a copy $K \in \mathcal{K}$: A vertex $v \in V$ is extending for a $K \in \mathcal{K}$ if K is contained in $\mathcal{H}_v$, or in other words, if for each edge e of K, $e \dot\cup \{v\}$ is an edge in $\mathcal{H}$. The picture below shows a $K^2(2)$ and an extending vertex for it (in a 3-regular hypergraph).

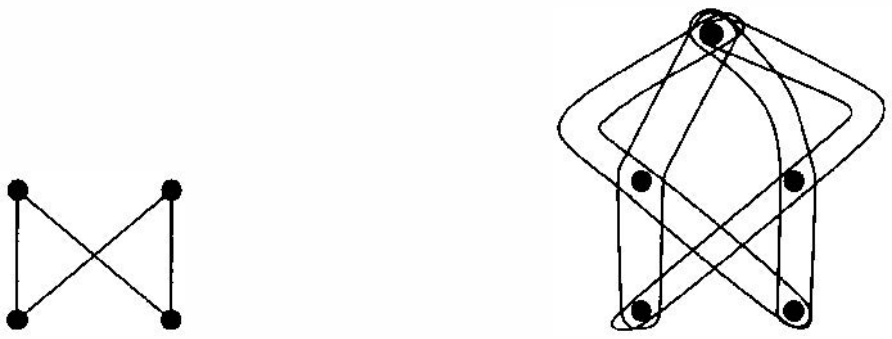

The idea is to count the number of all pairs (K, v), where $K \in \mathcal{K}$ and v is an extending vertex of K, in two ways.

On the one hand, if a fixed copy $K \in \mathcal{K}$ has q_K extending vertices, then it contributes $\binom{q_K}{t}$ distinct copies of $K^k(t)$ in $\mathcal{H}$. We note that one copy of $K^k(t)$ comes from at most $O(1)$ distinct $K \in \mathcal{K}$ in this way, and therefore it suffices to bound $\sum_{K\in\mathcal{K}} \binom{q_K}{t}$ from below.

On the other hand, for a fixed vertex v, the hypergraph $\mathcal{H}_v$ contains at least $f_{k-1}(n, m_v)$ copies $K \in \mathcal{K}$ by the inductive assumption, where m_v is the number of edges of $\mathcal{H}_v$. Hence

$$\sum_{K\in\mathcal{K}} q_K \geq \sum_{v\in V} f_{k-1}(n, m_v).$$

Using $\sum_{v\in V} m_v = km$, the convexity of f_{k-1} in the second variable, and Jensen's inequality (see page xvi), we obtain

$$\sum_{K\in\mathcal{K}} q_K \geq n\, f_{k-1}(n, km/n)\,. \tag{9.1}$$

To conclude the proof, we define a convex function extending the binomial coefficient $\binom{x}{t}$ to the domain $\mathbf{R}$:

$$g(x) = \begin{cases} 0 & \text{for } x \leq t-1, \\ \frac{x(x-1)\cdots(x-t+1)}{t!} & \text{for } x > t-1\,. \end{cases}$$

We want to bound $\sum_{K\in\mathcal{K}} g(q_K)$ from below, and we have the bound (9.1) for $\sum_{K\in\mathcal{K}} q_K$. Using the bound $|\mathcal{K}| \le n^{t(k-1)}$ (clear, since $K^{k-1}(t)$ has $t(k-1)$ vertices) and Jensen's inequality, we derive that the number of copies of $K^k(t)$ in $\mathcal{H}$ is at least

$$cn^{t(k-1)} g\left(\frac{n f_{k-1}(n, km/n)}{n^{t(k-1)}}\right).$$

A calculation finishes the induction step; we omit the details. □

Bibliography and remarks. The second selection lemma was conjectured, and proved in the planar case, by Bárány, Füredi, and Lovász [BFL90]. The missing part for higher dimensions was the colored Tverberg theorem (discussed in Section 8.3). A proof for the planar case by a different technique, with considerably better quantitative bounds than can be obtained by the method shown above, was given by Aronov, Chazelle, Edelsbrunner, Guibas, Sharir, and Wenger [ACE+91] (the bounds were mentioned in the text). The full proof of the second selection lemma for arbitrary dimension appears in Alon, Bárány, Füredi, and Kleitman [ABFK92].

Several other "selection lemmas," sometimes involving geometric objects other than simplices, were proved by Chazelle, Edelsbrunner, Guibas, Herschberger, Seidel, and Sharir [CEG+94].

Theorem 9.2.2 is from Erdős and Simonovits [ES83].

Exercises

1. (a) Prove a one-dimensional selection lemma: Given an n-point set $X \subset \mathbf{R}$ and a family $\mathcal{F}$ of $\alpha\binom{n}{2}$ X-intervals, there exists a point common to $\Omega(\alpha^2\binom{n}{2})$ intervals of $\mathcal{F}$. What is the best value of the constant of proportionality you can get? [3]
   (b) Show that this result is sharp (up to the value of the multiplicative constant) in the full range of α. [2]
2. (a) Show that the exponent s_2 in the second selection lemma in the plane cannot be smaller than 2. [2]
   (b) Show that $s_3 \ge 2$. [4] Can you also show that $s_d \ge 2$?
   (c) Show that the proof method via the fractional Helly theorem cannot give a better value of s_2 than 3 in Theorem 9.2.1. That is, construct an n-point set and $\alpha\binom{n}{3}$ triangles on it in such a way that no more than $O(\alpha^5 n^9)$ triples of these triangles have a point in common. [2]

9.3 Order Types and the Same-Type Lemma

The order type of a set. There are infinitely many 4-point sets in the plane in general position, but there are only two "combinatorially distinct" types of such sets: