Topological Methods in Combinatorics

László Lovász

May 14, 2016

Contents

1 Introduction

Trivial: If G is connected and $|V| > 1$, then V has a partition $V = V_1 \cup V_2$ such that $G[V_i]$ is connected. It has another partition such that $G[V_1, V_2]$ is connected. It has a connected subgraph of every size $1 \leq k \leq n$.

Can we combine connectivity conditions and size conditions? Yes if we strengthen the connectivity assumption on G.

Proposition 1.1 *Let $G = (V, E)$ be a 2-connected graph and $1 \leq k \leq n$. Then V has a partition $V = V_1 \cup V_2$ such that $G[V_i]$ is connected and $|V_1| = k$. If n is even and G is non-bipartite, then G has another partition such that $G[V_1, V_2]$ is connected and $|V_1| = |V_2|$.*

Proof via a discrete version of Bolzano's Theorem and the following lemma.

Lemma 1.2 *Let $G = (V, E)$ be a 2-connected graph and let T_1, T_2 be two spanning trees. The T_1 can be transformed into T_2 by changing a single edge incident with a leaf at each step. We can fix a root that is not considered as a leaf.*

What happens if we need a partition $V = V_1 \cup V_2 \cup V_3$ into sets with prescribed sizes? We need more involved topology than Bolzano's Theorem.

2 Combinatorial homotopy theory

2.1 Preliminaries

A *simplicial complex* $\mathcal{K}$ is a finite collection of nonempty finite sets such that $X \in \mathcal{K}$, $Y \subseteq X \neq \emptyset$ implies $Y \in \mathcal{K}$. The union of all members of $\mathcal{K}$ is denoted by $V(\mathcal{K})$. The elements of $V(\mathcal{K})$ are called the *vertices* of $\mathcal{K}$, the elements of $\mathcal{K}$ are called the *simplices* of $\mathcal{K}$.

The *dimension* of a simplex $S \in \mathcal{K}$ is $\dim(S) = |S| - 1$. The *dimension* of $\mathcal{K}$ is the maximum dimension of any simplex in $\mathcal{K}$. The *k-dimensional skeleton* of a simplicial complex $\mathcal{K}$ is the simplicial complex $\mathcal{K}|_k = \{S \in \mathcal{K} : \ |X| \leq k+1\}$. We define the *link* of a vertex $v \in V(\mathcal{K})$ as the simplicial complex

$$\mathrm{lk}_{\mathcal{K}}(v) = \{X \subseteq V(K) \setminus \{v\} : \ X \cup \{v\} \in \mathcal{K}\}.$$

Let $n = |V(\mathcal{K})|$. Embed $V(\mathcal{K})$ in $\mathbb{R}^{n-1}$ so that the n points representing the vertices are not contained in one hyperplane. For each simplex S, take the convex hull $\mathrm{conv}(S)$ of the elements in S, and take the union of these simplices. This set $G(\mathcal{K}) \subset \mathbb{R}^{n-1}$ is called a *geometric realization* of $\mathcal{K}$. The sets $\mathrm{conv}(S)$ $(S \in \mathcal{K})$ are the *faces* of $G(\mathcal{K})$. By the "general position" assumption on the vertices, it follows that the faces intersect each other in their common subface only: $\mathrm{conv}(S_1) \cap \mathrm{conv}(S_2) = \mathrm{conv}(S_1 \cap S_2)$.

If two topological spaces T_1, T_2 are homeomorphic, then we write $T_1 \cong T_2$. It is clear that all geometric realizations of the same simplicial complex are homeomorphic. We say that two simplicial complexes are homeomorphic, if their geometric realization are. (We use similar terminology for other topological properties of geometric realizations of simplicial complexes.)

Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be two simplicial complexes and let $f: V(\mathcal{K}_1) \to V(\mathcal{K}_2)$ be a mapping. We say that f is *simplicial* if $f(S) \in \mathcal{K}_2$ for every simplex $S \in \mathcal{K}_1$. Every simplicial mapping can be extended linearly to get a continuous mapping $\widehat{f}: G(\mathcal{K}_1) \to G(\mathcal{K}_2)$.

The *baricentric subdivision* $\mathcal{B}(\mathcal{K})$ of a simplicial complex $\mathcal{K}$ is obtained as follows. We create a new vertex v_S for every non-empty simplex S (the "baricenter" of S), and define $v_{S_1}, \ldots, v_{S_k}$ to be a simplex for every $S_1 \subset \cdots \subset S_k$. It is easy to see that $G(\mathcal{B}(\mathcal{K}))$ is homeomorphic with $G(\mathcal{K})$. In fact, there is a "canonical" homeomorphism $\beta: G(\mathcal{B}(\mathcal{K})) \to G(\mathcal{K})$, where $\beta(v_S)$ is the baricenter of $\mathrm{conv}(S)$ (which is a face of $G(\mathcal{K})$), and β is extended linearly over each face in $G(\mathcal{B}(\mathcal{K}))$.

Example 2.1 The *simplex* on a finite set V is the simplicial complex $\Sigma(V) = 2^V \setminus \{\emptyset\}$. The *boundary* of the simplex $\Sigma(V)$ is the simplicial complex $\Gamma(V) = 2^V \setminus \{\emptyset, V\}$. The usual geometric representation of $\Sigma(V)$ is the standard simplex $\Delta^V = \{x \in \mathbb{R}^V : x_i \geq 0, \sum_i x_i \leq 1\}$. If $|V| = \{0, 1, \ldots, r\}$, then we denote $\Sigma(V)$ by Σ^r and $\Gamma(V)$ by Γ^{r-1}.

The r-dimensional closed unit ball is denoted by B^r. Specifically, B^0 is a single point and B^1 is a segment. The boundary of B^r is the $(r-1)$-dimensional unit sphere, denoted by S^{r-1}. So S^0 consists of two points, and S^1 is a circle. Sometimes it is convenient to define $S^{-1} = \emptyset$. Clearly,

$$G(\Sigma^r) \cong B^r \quad \text{and} \quad G(\Gamma^{r-1}) \cong S^{r-1}. \tag{1}$$

♦

2.2 Homotopy and homotopy equivalence

Let T_1 and T_2 be two topological spaces. Two continuous maps $f_0, f_1 : T_1 \to T_2$ are called *homotopic* (denoted $f_0 \sim f_1$), if they can be deformed into each other. More exactly, there exists a continuous mapping $F: T_1 \times [0,1] \to T_2$ such that $f_0(x) = F(x, 0)$ and $f_1(x) = F(x, 1)$.

Homotopy of maps is an equivalence relation. In particular, if $f_1, f_2, f_3 : T_1 \to T_2$ and $f_1 \sim f_2$ and $f_2 \sim f_3$, then $f_1 \sim f_3$. If $f: T_1 \to T_2$, $g_1, g_2 : T_2 \to T_3$, $h: T_3 \to T_4$, and $g_1 \sim g_2$, then $f \circ g_1 \circ h \sim f \circ g_2 \circ h$.

Example 2.2 Let $\mathrm{const}_p : T_1 \to T_2$ denote the map which maps every point to $p \in T_2$. For every convex set K and point $p \in K$, the maps $\mathrm{id}_K, \mathrm{const}_p : K \to K$ are homotopic. ♦

We say that T_1 and T_2 are *homotopy equivalent* (denoted by $T_1 \sim T_2$), if there exist continuous maps $f: T_1 \to T_2$ and $g: T_2 \to T_1$ such that $f \circ g \sim \mathrm{id}_{T_2}$ and $g \circ f \sim \mathrm{id}_{T_1}$.

Homotopy equivalence is an equivalence relation. The equivalence classes of this relation are often called *homotopy types.*

Example 2.3 Every convex set is homotopy equivalent to a point. An annulus is homotopy equivalent to a circle. The union of two disjoint disks is not homotopy equivalent to a single disk. A circle is not homotopy equivalent with the one-point space. ⧫

Let $T_1 \subseteq T_2$ be two topological spaces. We say that T_1 is a *retract* of T_2 if there is a continuous map $\varphi: T_2 \to T_1$ such that $\varphi|_{T_1} = \mathrm{id}_{T_1}$. If there is a homotopy $F: T_2 \times [0,1]$ such that $F(x,t) = x$ for all $x \in T_1$ and $t \in [0,1]$, $F(x,0) = x$ for all $x \in T_2$ and $F(x,1) = \varphi(x)$ for all $x \in T_2$, then we call T_1 a *deformation retract* of T_2.

Every deformation retract of a space is homotopy equivalent to it. This does not remain valid for retracts in general.

Example 2.4 Any face of a convex polytope is a retract of the polytope. The 2-point space consisting of the endpoints of a segment is not a retract of this segment. An arc of a circle is a retract, but not a deformation retract, of the full circle. ⧫

Example 2.5 The following example will be useful later on. Let $A_1, \dots, A_m$ be finite sets, $r_1, \dots, r_k$ positive integers, let $V = A_1 \cup \dots \cup A_m$, and let $\mathcal{M}(A_1, \dots, A_m)$ consist of those subsets $X \subseteq V$ for which $|X \cap A_i| \le r_i$ for every $1 \le i \le m$. Then $\mathcal{M}(A_1, \dots, A_m)$ is a simplicial complex.

An interesting special case is obtained when the sets $A_1, \dots, A_m$ are disjoint and have two elements. In this case the simplices of $\mathcal{M}(A_1, \dots, A_m)$ can be identified with the proper faces of the m-dimensional cross-polytope $X_n = \{x \in \mathbb{R}^n : \sum_i |x_i| \le 1\}$. So $\mathcal{M}(A_1, \dots, A_m) \cong S^{m-1}$.

Suppose that $A_1, \dots, A_m$ are disjoint, and let $A_i' \subseteq A_i$, then $\mathcal{M}(A_1', \dots, A_m')$ is a retract of $\mathcal{M}(A_1, \dots, A_m)$: we can take any map $\varphi_i: A_i \to A_i'$ that is the identity on A_i' ($i = 1, \dots, m$), then $\varphi = \varphi_1 \cup \dots \cup \varphi_m$ is a simplicial map $\mathcal{M}(A_1, \dots, A_m) \to \mathcal{M}(A_1', \dots, A_m')$ that is a retraction. ⧫

A topological space T is *contractible* if it is homotopy equivalent to the space with a single point. This is equivalent to saying that the map const_p (where $p \in T$) is homotopic to the identity map id_T. This property does not depend on the choice of p.

Any two continuous maps from a topological space into a contractible space are homotopic. Every convex body is contractible. A simple graph (as a 1-dimensional simplicial complex) is contractible if and only if it is a tree. If a simplicial complex has a vertex u that is contained in every maximal simplex, then it is contractible.

The bad news: given a simplicial complex, it is algorithmically undecidable whether it is contractible. This implies that it is algorithmically undecidable whether two simplicial complexes are homotopy equivalent, or two simplicial maps are homotopic. It is somewhat surprising that in spite of these facts, there is a useful combinatorial theory of homotopy equivalence and contractibility.

Lemma 2.6 *Let T be a topological space and let $\mathcal{K}$ be a simplicial complex. Let $f, g : T \to G(\mathcal{K})$ be two continuous maps such that $f(x)$ and $g(x)$ are contained in the same face of $G(\mathcal{K})$ for all $x \in T$. Then $f \sim g$.*

Proof. $\Phi(t, x) = (1-t)f(x) + tg(x) \in G(\mathcal{K}_2)$ defines a homotopy of f and g. □

Lemma 2.7 *Every retract of a contractible space is contractible.*

Proof. Let $T_1 \subseteq T_2$ be a topological spaces, and let $\varphi : T_1 \to T_2$ be a retraction. Assume that T_1 is contractible, let $p \in T_1$, and let $F : T_1 \times [0,1]$ be a homotopy such that $F(x, 0) = x$ and $F(x, 1) = p$ for $x \in T_1$. The $F'(x, t) = \varphi(F(x, t))$, considered as a map $T_2 \times [0,1] \to T_2$, is a homotopy with $F'(x, 0) = \varphi(F(x, 0)) = \varphi(x) = x$ and $F'(x, 1) = \varphi(F(x, 1)) = \varphi(p) = p$ for $x \in T_2$. □

Lemma 2.8 *If a topological space $\mathcal{T}$ is contractible, then every continuous map $f : S^{r-1} \to \mathcal{T}$ extends to a continuous map $\widehat{f} : B^r \to \mathcal{T}$.*

Proof. There is a homotopy $F : S^{r-1} \times [0,1]$ such that $F(., 0)$ is a constant map and $F(., 1)$ is the identity map. The map $\widehat{f}(x) = F(f(x/\|x\|), \|x\|)$ is a continuous map extending f. □

For simplicial complexes, we can prove a converse.

Lemma 2.9 *For a simplicial complex $\mathcal{K}$ the following are equivalent:*

(i) *$\mathcal{K}$ is contractible;*

(ii) *For every finite set V, every continuous map $\Gamma(V) \to G(\mathcal{K})$ extends to a continuous map $\Sigma(V) \to G(\mathcal{K})$.*

(iii) *If $\mathcal{K}_1 \subseteq \mathcal{K}_2$ are simplicial complexes, then every continuous map $G(\mathcal{K}_1) \to G(\mathcal{K})$ extends to a continuous map $G(\mathcal{K}_2) \to G(\mathcal{K})$.*

(iv) *$\mathcal{K}$ is a retract of $\Sigma(V(\mathcal{K}))$.*

Proof. (i)⇒(ii) by Lemma 2.8.

(ii)⇒(iii): Let S be a simplex in $\mathcal{K}_1 \setminus \mathcal{K}_2$ of minimal dimension. Then $\Gamma(S) \subseteq \mathcal{K}_1$ and so the map f is already defined on $G(\Gamma(S))$. Hence by (ii), f extends to a continuous map $G(\mathcal{K}_1 \cup \{S\}) \to G(\mathcal{K})$. Repeating this argument, we extend f to a continuous map $G(\mathcal{K}_2) \to G(\mathcal{K})$.

(iii)⇒(iv): The map $\mathrm{id}_{\mathcal{K}}$ defines a map $G(\mathcal{K}) \to G(\Sigma(V(\mathcal{K})))$, which by (iii) extends to a continuous map $G(\Sigma(V(\mathcal{K}))) \to G(\Sigma(V(\mathcal{K})))$, which is a retraction.

(iv)⇒(i): Since $\Sigma(V)$ is contractible, every retract of it is contractible as well. □

Lemma 2.10 *Let $\mathcal{K}$ be a simplicial complex and $U \subseteq V(\mathcal{K})$. Suppose that $\mathcal{K} \cap \Sigma(U)$ is contractible. Then $\mathcal{K} \cup \Sigma(U) \sim \mathcal{K}$.*

Proof. By Lemma 2.9, there is a retraction $\varphi_0 : G(\Sigma(U)) \to G(\mathcal{K} \cap \Sigma(U))$. Define $\varphi : G(\mathcal{K} \cup \Sigma(U)) \to G(\mathcal{K} \cup \Sigma(U))$ by

$$\varphi(x) = \begin{cases} \varphi_0(x) & \text{if } x \in G(\Sigma(U)), \\ x, & \text{otherwise.} \end{cases}$$

Clearly $\varphi(x) \in G(\mathcal{K})$ for every x, and φ, as a map $G(\mathcal{K} \cup \Sigma(U)) \to G(\mathcal{K})$, is a retraction. Since x and $\varphi(x)$ are contained in the same simplex of $\mathcal{K} \cup \Sigma(U)$, it follows that $\varphi \sim \mathrm{id}_{G(\mathcal{K} \cup \Sigma(U))}$. Hence φ is a deformation retraction, showing that $\mathcal{K} \cup \Sigma(U) \sim \mathcal{K}$. □

This lemma generalizes to adding more than one simplex.

Lemma 2.11 *Let $\mathcal{K}$ be a simplicial complex and let $U_1, \ldots, U_m \subseteq V(\mathcal{K})$. Suppose that for every $1 \le i_1 < \cdots < i_r \le m$, the restriction $\mathcal{K} \cap \Sigma(U_{i_1} \cap \cdots \cap U_{i_r})$ is either empty or contractible. Then $\mathcal{K} \sim \mathcal{K} \cup \Sigma(U_1) \cup \cdots \cup \Sigma(U_m)$.*

Proof. By induction on m. If $m = 0$, then we have nothing to prove. Let $m \ge 1$. By Lemma 2.10, we have $\mathcal{K} \sim \mathcal{K}' = \mathcal{K} \cup \Sigma(U_1)$. Furthermore,

$$\mathcal{K}' \cap \Sigma(U_{i_1} \cap \cdots \cap U_{i_r}) = (\mathcal{K} \cap \Sigma(U_{i_1} \cap \cdots \cap U_{i_r})) \cup \Sigma(U_1 \cap U_{i_1} \cap \cdots \cap U_{i_r})$$

is either empty or contractible for every $2 \le i_1 < \cdots < i_r \le m$, again by Lemma 2.10. By the induction hypothesis, this implies that

$$\mathcal{K}' \sim \mathcal{K}' \cup \Sigma(U_2) \cup \cdots \cup \Sigma(U_m) = \mathcal{K} \cup \Sigma(U_1) \cup \cdots \cup \Sigma(U_m).$$

This proves the Lemma. □

Lemma 2.12 *Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be simplicial complexes. If $\mathcal{K}_1$, $\mathcal{K}_2$ and $\mathcal{K}_1 \cap \mathcal{K}_2$ are contractible, then so is $\mathcal{K}_1 \cup \mathcal{K}_2$.*

Proof. Let $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2$, $V_i = V(\mathcal{K}_i)$, $V = V_1 \cup V_2$ and $S = V_1 \cap V_2$. The hypothesis that $\mathcal{K}_1 \cap \mathcal{K}_2$ is contractible implies that S is not empty. The proof goes in several steps.

1°. Assume that $\mathcal{K}_1 \cap \mathcal{K}_2 = \Sigma(S)$. Since $\mathcal{K}_i$ is contractible, Lemma 2.9 implies that there are retractions $\varphi_i : G(\Sigma(V_i)) \to G(\mathcal{K}_i)$. Since $\varphi_i|_G(S) = \mathrm{id}$, we can "glue together" φ_1 and φ_2 to get $\varphi = \varphi_1 \cup \varphi_2$, to get a retraction from $G(\Sigma(V_1) \cup \Sigma(V_2)) \to G(\mathcal{K})$. Since $\Sigma(V_1) \cup \Sigma(V_2)$ is trivially contractible (any point of S is contained in both maximal simplices), it follows by Lemma 2.7 that so is $\mathcal{K}$.

2°. Next, we treat the more general (but still not completely general) case when $\mathcal{K}_1 \cap \Sigma(S) = \mathcal{K}_2 \cap \Sigma(S)$. Then $\mathcal{K} \cap \Sigma(S) = \mathcal{K}_i \cap \Sigma(S) = \mathcal{K}_1 \cap \mathcal{K}_2$ is contractible, and hence by Lemma 2.10, we have $\mathcal{K} \sim \mathcal{K} \cup \Sigma(S) = (\mathcal{K}_1 \cup \Sigma(S)) \cup (\mathcal{K}_2 \cup \Sigma(S))$. By Lemma 2.10 again, $\mathcal{K}_1 \cup \Sigma(S)$ and $\mathcal{K}_2 \cup \Sigma(S)$ are contractible, hence by the first part of the proof, so is $(\mathcal{K}_1 \cup \Sigma(S)) \cup (\mathcal{K}_2 \cup \Sigma(S))$, and hence so is $\mathcal{K}$.

3°. Finally, in the general case, consider the baricentric subdivisions $\mathcal{B}(\mathcal{K}_i)$. These are contractible, and so is $\mathcal{B}(\mathcal{K}_1) \cap \mathcal{B}(\mathcal{K}_2) = \mathcal{B}(\mathcal{K}_1 \cap \mathcal{K}_2)$. Furthermore, we have $V(\mathcal{B}(\mathcal{K}_1)) \cap V(\mathcal{B}(\mathcal{K}_2)) = \mathcal{K}_1 \cap \mathcal{K}_2$, and the restriction of $\mathcal{B}(\mathcal{K}_i)$ to this is $\mathcal{B}(\mathcal{K}_1 \cap \mathcal{K}_2)$ for $i = 1, 2$. Hence the previous special case applies, and we get that $\mathcal{B}(\mathcal{K}_1) \cup \mathcal{B}(\mathcal{K}_2) = \mathcal{B}(\mathcal{K}_1 \cup \mathcal{K}_2)$ is contractible. This implies that $\mathcal{K}_1 \cup \mathcal{K}_2$ is contractible. □

Corollary 2.13 *Let $\mathcal{K}$ be a simplicial complex and $v \in V(\mathcal{K})$. Suppose that both $\mathcal{K} \setminus v$ and $\mathrm{lk}_{\mathcal{K}}(v)$ are contractible. Then $\mathcal{K}$ is contractible.*

Let $\mathcal{K}$ be a simplicial complex, and let $f: U \to V(\mathcal{K})$ be any mapping. Define the simplicial complex

$$f^{-1}(\mathcal{K}) = \{X \subseteq U : f(X) \in \mathcal{K}\}.$$

Lemma 2.14 (Contractible Carrier Lemma) *Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be simplicial complexes and let $f: V(\mathcal{K}_1) \to V(\mathcal{K}_2)$ be a surjective simplicial map. Suppose that for every simplex $S \in \mathcal{K}_2$, the complex $\mathcal{K}_1 \cap \Sigma(f^{-1}(S))$ is contractible. Then $\mathcal{K}_1 \sim \mathcal{K}_2$.*

The proof will show that it would suffice to require that $\mathcal{K}_1 \cap \Sigma(f^{-1}(S))$ is contractible for every simplex S that is the intersection of maximal simplices.

Proof. We start with proving the special case when $\mathcal{K}_1 = f^{-1}(\mathcal{K}_2)$. In this case $f^{-1}(S) \in \mathcal{K}_1$ for every $S \in \mathcal{K}_2$, so the contractibility condition is trivially satisfied. By the definition of $f^{-1}(\mathcal{K}_2)$, the mapping $f: V(\mathcal{K}_1) \to V(\mathcal{K}_2)$ is simplicial, and so it extends to a continuous mapping $\widehat{f}: G(\mathcal{K}_1) \to G(\mathcal{K}_2)$.

For every $u \in V(\mathcal{K}_2)$, let $g(u) \in G(\mathcal{K}_1)$ be the center of gravity of the simplex $f^{-1}(u) \in \mathcal{K}_1$. If $S = \{u_1, \ldots, u_d\} \in \mathcal{K}_2$, then $f^{-1}(S) \in \mathcal{K}_1$, and hence $g(u_1), \ldots, g(u_d)$ are contained in the face $\mathrm{conv}(f^{-1}(S))$ of $G(\mathcal{K}_1)$. Hence we can extend the map g linearly to a continuous map $\widehat{g}: G(\mathcal{K}_2) \to G(\mathcal{K}_1)$.

It is clear that $\widehat{g} \circ \widehat{f} = \mathrm{id}_{G(\mathcal{K}_2)}$. On the other hand, let $x \in G(\mathcal{K}_1)$, say $x \in \mathrm{conv}(S)$, where $S \in \mathcal{K}_1$. Then both x and $\widehat{f}(\widehat{g}(x))$ are contained in the face $\mathrm{conv}(f^{-1}(f(S))$, and hence $\widehat{f} \circ \widehat{g} \sim \mathrm{id}_{G(\mathcal{K}_1)}$ by Lemma 2.6. This proves that $\mathcal{K}_1 \sim \mathcal{K}_2$.

Now in the general case, we already know that $f^{-1}(\mathcal{K}_2) \sim \mathcal{K}_2$. Let $S_1, \ldots, S_m$ be the maximal simplices in $\mathcal{K}_2$, and let $U_i = f^{-1}(S_i)$. Then

$$\mathcal{K}_1 \cap \Sigma(U_{i_1} \cap \cdots \cap U_{i_r}) = \mathcal{K}_1 \cap \Sigma(f^{-1}(S_{i_1} \cap \cdots \cap S_{i_r}))$$

is contractible for every $1 \leq i_1 < \cdots < i_r$, and hence by Lemma 2.11,

$$\mathcal{K}_1 \sim \mathcal{K}_1 \cup \Sigma(U_1) \cup \cdots \cup \Sigma(U_m) = \Sigma(U_1) \cup \cdots \cup \Sigma(U_m) = f^{-1}(\mathcal{K}_2)$$

Thus $\mathcal{K}_1 \sim f^{-1}(\mathcal{K}_2) \sim \mathcal{K}_2$. $\square$

2.3 The Nerve Theorem

Let $G = (V, E)$ be a bipartite graph with bipartition $V = U \cup W$. The *neighborhood complex* $\mathcal{N}_U = \mathcal{N}_U(G)$ of G is the simplicial complex consisting of all subsets A of U such that the elements of A have a common neighbor. The neighborhood complex $\mathcal{N}_W = \mathcal{N}_W(G)$ is defined analogously.

Lemma 2.15 *The two neighborhood complexes of a bipartite graph are homotopy equivalent.*

Proof. For every simplex $S \in \mathcal{N}_U$, let $f(S) \subseteq W$ denote the set common neighbors of nodes in S. Clearly $f(S) \in \mathcal{N}_W$, so $f:\ \mathcal{N}_U \to \mathcal{N}_W$. We define $g:\ \mathcal{N}_W \to \mathcal{N}_U$ analogously.

If $S_1 \subset \cdots \subset S_k$, then $f(S_1) \supset \cdots \supset f(S_k)$, and hence f is a simplicial map between the baricentric subdivisions: $f:\ V(\mathcal{B}(\mathcal{N}_U)) = \mathcal{N}_U \to V(\mathcal{B}(\mathcal{N}_W)) = \mathcal{N}_W$. We define $g:\ \mathcal{N}_W \to \mathcal{N}_U$ analogously. We claim that f and g certify the homotopy equivalence of $\mathcal{B}(\mathcal{N}_U)$ and $\mathcal{B}(\mathcal{N}_W)$.

The composite map $h = f \circ g:\ \mathcal{N}_U \to \mathcal{N}_U$ is a simplicial map $\mathcal{B}(\mathcal{N}_U) \to \mathcal{B}(\mathcal{N}_U)$, and so it defines a continuous map $\widehat{h}:\ G(\mathcal{B}(\mathcal{N}_U)) \to G(\mathcal{B}(\mathcal{N}_U))$. There is a canonical map $\alpha:\ G(\mathcal{B}(\mathcal{N}_U)) \to G(\mathcal{N}_U)$, and so $\widehat{h} \circ \alpha:\ G(\mathcal{B}(\mathcal{N}_U)) \to G(\mathcal{N}_U)$ is a continuous map.

The map h has the property that $h(S) \supseteq S$ for every $S \in \mathcal{N}_U$. Hence if $S_1 \subset \cdots \subset S_k$ is a simplex of $\mathcal{B}(\mathcal{N}_U)$, then $S_1, \ldots, S_k$ as well as $h(S_1), \ldots, h(S_k)$ are subsets of $h(S_k)$. This implies that for every $x \in G(\mathcal{B}(\mathcal{N}_U))$, the points $\widehat{h} \circ \alpha(x)$ and $\alpha(x)$ are both contained in the same face $G(h(S_k))$. By Lemma 2.6, this implies that $\widehat{h} \circ \alpha \sim \alpha$, and hence $\widehat{h} = (\widehat{h} \circ \alpha) \circ \alpha^{-1} \sim \alpha \circ \alpha^{-1} = \mathrm{id}$.

We argue similarly about $g \circ f$, and so $\mathcal{B}(\mathcal{N}_U)) \sim \mathcal{B}(\mathcal{N}_W)$, which is equivalent to $\mathcal{N}_U \sim \mathcal{N}_W$. □

Let $\mathcal{H}$ be a hypergraph consisting of nonempty sets. The *nerve* of $\mathcal{H}$ is defined as the simplicial complex $\mathsf{nerve}(\mathcal{H})$ whose vertices are the sets in $\mathcal{H}$, and $\{X_1, \ldots, X_r\} \in \mathsf{nerve}(\mathcal{H})$ if and only if $X_1, \ldots, X_r \in \mathcal{H}$ and $X_1 \cap \cdots \cap X_r \neq \emptyset$.

Corollary 2.16 *Let $\mathcal{K}$ be a simplicial complex, and assume that $\mathcal{H} \subseteq \mathcal{K}$ contains all maximal simplices in $\mathcal{K}$. Then $\mathcal{K} \sim \mathsf{nerve}(\mathcal{H})$. In particular, $\mathcal{K} \sim \mathsf{nerve}(\mathcal{K})$.*

More generally:

Theorem 2.17 (Nerve Theorem) *Let $\mathcal{K}$ be a simplicial complex and let $\mathcal{K}_1, \ldots, \mathcal{K}_m$ be subcomplexes such that $\mathcal{K} = \mathcal{K}_1 \cup \cdots \cup \mathcal{K}_m$. Assume that $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r}$ is either empty or contractible for every $1 \le i_1 < \cdots < i_r \le m$, $r \ge 1$. Then $\mathcal{K} \sim \mathsf{nerve}\{\mathcal{K}_1, \ldots, \mathcal{K}_m\}$.*

Proof. Let $V_i = V(\mathcal{K}_i)$. First, we prove the case when $\mathcal{K}_i = \mathcal{K} \cap \Sigma(V_i)$. In this case $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r} = \mathcal{K} \cap \Sigma(V_{i_1} \cap \cdots \cap V_{i_r})$. Lemma 2.11 implies that $\mathcal{K} \sim \mathcal{K} \cup \Sigma(V_1) \cup \cdots \cup \Sigma(V_m) = \Sigma(V_1) \cup \cdots \cup \Sigma(V_m)$. By Corollary 2.16, we have $\Sigma(V_1) \cup \cdots \cup \Sigma(V_m) \sim \mathsf{nerve}\{V_1, \ldots, V_m\}$, which proves the theorem in this case.

For the general case (when $\mathcal{K}_i$ may be smaller than $\mathcal{K} \cap \Sigma(V_i)$), we replace each $\mathcal{K}_i$ by $\mathcal{B}(\mathcal{K}_i)$ as in the proof of Lemma 2.12. □

To formulate another important consequence of the Nerve Theorem, we need a construction of simplicial complexes generalizing baricentric subdivisions. Let $P = (V, \le)$ be a poset (partially ordered set), and let $\mathcal{C}(P)$ be the set of all nonempty chains (totally ordered subsets) of P. Clearly $\mathcal{C}(P)$ is a simplicial complex, called the *chain complex* of P.

Example 2.18 If $\mathcal{K}$ is a simplicial complex and $P = (\mathcal{K}, \subseteq)$, then $\mathcal{C}(P)$ is just the baricentric subdivision of $\mathcal{K}$. ♦

A *cross-cut* in a poset P is an antichain that meets every maximal chain. For example, if P is the set of all subsets of a finite set ordered by inclusion, then all sets of a given size form a crosscut. The following theorem, which concerns the case when P is obtained from a finite lattice, can be considered as a further version of the Nerve Theorem.

Theorem 2.19 (Cross-Cut Theorem) *Let L be a finite lattice with operations $\wedge$ and $\vee$, let $P = L \setminus \{0, 1\}$, and let S be a cross-cut in P. Then the simplicial complex*

$$\mathcal{A}_S = \{X \subseteq S \ : \ \vee X \neq 1 \text{ or } \wedge X \neq 0\}$$

is homotopy equivalent to $\mathcal{C}(P)$.

Proof. For $s \in S$, let $\mathcal{B}_s \subseteq \mathcal{C}(P)$ denote the set of those chains $X \subseteq P$ that contain s, or are contained in such a chain. The assumption that S is a cross-cut implies that $\mathcal{C}(P) = \cup_{s \in S} \mathcal{B}_s$. For $T \subseteq S$, let $\mathcal{B}_T = \cap_{s \in T} \mathcal{B}_s$.

Claim 2.20 *$\mathcal{B}_T \neq \emptyset$ if and only if $T \in \mathcal{A}_S$.*

Suppose that $\mathcal{B}_T \neq \emptyset$, and let $x \in \mathcal{B}_T$. Then x is comparable with every $s \in T$, and since S is an antichain, either $x \geq s$ for every $s \in T$, or $x \leq s$ for every $s \in T$. Hence both $\wedge T$ and $\vee T$ are comparable with x, and either $\wedge T \neq 0$ or $\vee T \neq 1$, implying that $T \in \mathcal{A}_S$. The converse follows similarly.

This argument also gives that if $\mathcal{B}_T \neq \emptyset$, then it has a vertex contained in every maximal simplex, and hence it is contractible. Thus the Nerve Theorem 2.17 applies, and gives that $\mathcal{C}(P) = \cup_{s \in S} \mathcal{B}_s \sim \mathsf{nerve}\{\mathcal{B}_s : \ s \in S\} \cong \mathcal{A}_S$. $\square$

Corollary 2.21 *Let L be a finite lattice, $P = L \setminus \{0, 1\}$, and let A be the set of atoms of L. Then the simplicial complex*

$$\mathcal{A} = \{X \subseteq A \ : \vee X \neq 1\}$$

is homotopy equivalent to $\mathcal{C}(P)$.

The complex $\mathcal{A}_S$ can be defined for a set $S \subseteq P$ in an arbitrary poset P: it consists of those subsets X of S that are bounded from either below or above in P. One might think that this leads to a generalization of the Cross-Cut Theorem. However, the poset obtained from a 2-element chain by doubling every element shows that the theorem is not true for posets in general.

3 Brouwer's fixed point theorem

A simplicial complex $\mathcal{K}$ such that $G(\mathcal{K}) \cong S^r$ is called a *triangulation* of the sphere S^r. For a coloring of its vertices, a simplex is called *colorful* if all its vertices get different colors.

Lemma 3.1 *Let $\mathcal{K}$ be a simplicial complex homeomorphic to the sphere S^r. In every coloring of the vertices of $\mathcal{K}$ with $r + 1$ colors, the number of r-dimensional colorful simplices is even.*

Proof. Let us count the number of colorful $(r-1)$-simplices with colors $1,\dots,r$. Every colorful r-simplex contains one such simplex. If a non-colorful r-simplex contains colorful $(r-1)$-simplex with colors $1,\dots,r$, then it contains exactly two such simplices. Since every colorful $(r-1)$-simplex with colors $1,\dots,r$ is counted exactly twice (for the two r-simplices on both sides), we must get an even number. Hence the number of colorful r-simplices must be even. $\square$

This lemma is often formulated differently. Let Δ^r denote the r-dimensional regular simplex with edge length 1, and let $\Phi_0,\dots,\Phi_r$ be its facets. By a *triangulation* of Δ^r we mean a family $\mathcal{F}$ of r-dimensional simplices such that $\cup\mathcal{F}=\Delta^r$ and $A\cap B$ is a face of both A and B for any two $A,B\in\mathcal{F}$. Such a triangulation defines a simplicial complex $\mathcal{K}$.

Lemma 3.2 (Sperner's lemma) *Let $\mathcal{K}$ be a triangulation of Δ^r. Let us color the vertices of $\mathcal{K}$ with $r+1$ colors $\{0,1,\dots,r\}$ so that vertices on face Φ_i are never colored with color i. Then there must be a colorful r-simplex in $\mathcal{K}$.*

Proof. We prove by induction on r that the number of colorful r-simplices is odd (and hence nonzero). Create a new vertex v, and add to $\mathcal{K}$ all simplices consisting of v and a simplex of K on one of the faces Φ_i. This way we get a simplicial complex $\mathcal{K}'$ homeomorphic to S^r. Let us color v with 0. By Lemma 3.1, we will have an even number of colorful r-simplices. Those colorful r-simplices that contain v must contain a colorful $(r-1)$-simplex from face F_0; the number of these is odd by induction. Hence the number of colorful r-simplices not containing v is odd as well. These are exactly the colorful r-simplices in $\mathcal{K}$. $\square$

The following theorem is an infinite version of Sperner's Lemma.

Theorem 3.3 *Let $A_0,\dots,A_r\subseteq\Delta^r$ be closed sets such that $A_0\cup\dots\cup A_r=\Delta^r$ and $A_i\cap\Phi_i=\emptyset$. Then $A_0\cap\dots\cap A_r\neq\emptyset$.*

Proof. Assume that $A_0\cap\dots\cap A_r=\emptyset$. Let $d(x,A_i)$ denote the distance of point x from A_i, then for every point $x\in\Delta^r$, we must have $\max(d(x,A_0),\dots,d(x,A_r))>0$. Since the left side is a continuous function of x, there is an $\varepsilon>0$ such that $\max(d(x,A_0),\dots,d(x,A_r))>\varepsilon$ for every x.

Let $\mathcal{K}$ be a triangulation of Δ^r such that the diameter of simplices in $\mathcal{K}$ is at most $\varepsilon/(d+1)$. Color each vertex u by one of the numbers i for which $u\in A_i$. Then the conditions of Sperner's Lemma are met, and hence there is a colorful r-simplex B. But then for any point $x\in B$, we have $\max(d(x,A_0),\dots,d(x,A_r))\le\varepsilon$, a contradiction. $\square$

Theorem 3.4 (Brouwer's fixed point theorem) *Every continuous map $B^r\to B^r$ has a fixed point.*

Proof. We may replace B^r by Δ^r, since they are homeomorphic. Let $f:\Delta^r\to\Delta^r$ be a continuous map, and suppose it has no fixed point, then there is an $\varepsilon>0$ such that $|f(x)-x|>\varepsilon$ for every $x\in\Delta^r$. Let e_i denote the outward oriented normal vector of unit length of the facet Φ_i. Let $A_i=\{x\in\Delta^r:\ (f(x)-x)\cdot e_i\ge\varepsilon/r\}$. Then:

— A_i is closed (trivially).

— $A_i \cap \Phi_i = \emptyset$, since $|f(x) - x| > 0$ (as f has no fixed point), but $(f(x) - x) \cdot e_i \le 0$ for $x \in \Phi_i$.

— $A_0 \cup \cdots \cup A_r = \Delta$. Let $x \in \Delta^r$. The inequalities $y \cdot e_i \le |f(x) - x|/r$ define a regular simplex centered at the origin with an inscribed ball of radius $|f(x) - x|/r$, and hence this simplex is contained in a ball with radius $|f(x) - x|$, which means that the point $y = f(x) - x$ is either outside this simplex or on its boundary. Thus there is an i for which $(f(x) - x) \cdot e_i \ge |f(x) - x|/r$, i.e., $x \in A_i$.

— $A_0 \cap \cdots \cap A_r = \emptyset$, since $\sum_i e_i = 0$, hence $\sum_i (f(x) - x) \cdot e_i = 0$ for every $x \in \Delta^r$, and so there must be an i for which $(f(x) - x) \cdot e_i \le 0$, and so $x \notin A_i$.

These four observations contradict Theorem 3.3. □

There are a number of equivalent assertions.

Corollary 3.5 (a) *If $\mathcal{K}$ is a contractible simplicial complex, then every continuous map $G(\mathcal{K}) \to G(\mathcal{K})$ has a fixed point.*

(b) *S^{r-1} is not a retract of B^r.*

(c) *Every continuous map $f: \Delta^r \to \Delta^r$ such that $f(\Phi_i) \subseteq \Phi_i$ for every $0 \le i \le r$ is surjective.*

Proof. (a) Let $f: G(\mathcal{K}) \to G(\mathcal{K})$ be a continuous map. By Lemma 2.9, there is a retraction $\varphi: G(\Sigma(V(\mathcal{K}))) \to G(\mathcal{K})$. The map $\varphi \circ f: G(\Sigma(V(\mathcal{K}))) \to G(\Sigma(V(\mathcal{K})))$ has a fixed point x by Theorem 3.4. Clearly $x \in G(\mathcal{K})$, and $x = f(\varphi(x)) = f(x)$.

(b) Assume that there is a retraction $\varphi: B^r \to S^{r-1}$, then $x \mapsto -\varphi(x)$ is a map $B^r \to B^r$ with no fixed point, a contradiction.

(c) Suppose that $a \in \Delta^r$ is not in the range of f. Since $f(\Delta^r)$ is closed, there is a neighborhood of a disjoint from $f(\Delta^r)$, and so we may assume that a is an internal point of Δ^r. For $x \in \Delta^r$, let $g(x)$ be the point where the semiline from $f(x)$ through a meets the boundary of Δ^r. Since $a \notin f(\Delta^r)$, the function $g: \Delta^r \to \Delta^r$ is well defined and continuous, and hence it must have a fixed point y by Theorem 3.4. Since $y = g(y)$, it must be on the boundary of Δ^r, i.e., $y \in \Phi_i$ for some $0 \le i \le r$. By hypothesis, $f(y) \in \Phi_i$, which contradicts the definition of g. □

We also need the following group-theoretic corollary of Brouwer's Fixed Point Theorem.

Corollary 3.6 *Let T be a compact contractible topological space and let G be a finite cyclic group acting on T. Then the elements of G have a common fixed point.*

Proof. Let φ be the action of a generator of G. Then φ has a fixed point by Corollary 3.5, and this is a fixed point of every other element of G as well. □

This last corollary does not remain valid for all finite groups, but it does remain valid for certain non-cyclic finite groups. One such class of finite groups is the following.

Theorem 3.7 [43] *Let T be a compact contractible topological space and let Γ be a finite group acting on T. Assume that Γ has a normal p-subgroup Γ_1 such that Γ/Γ_1 is cyclic. Then the elements of G have a common fixed point.* □

Example 3.8 To illustrate how to apply these results, let us return to Example 2.5. We show that if $|A_i| \geq 2$ for all $1 \leq i \leq m$, then $\mathcal{M}(A_1, \dots, A_m)$ is not $(m-1)$-connected. Let $A'_i = \{u_i, v_i\} \subseteq A_i$. The embedding $\mathcal{M}(A'_1, \dots, A'_m) \to \mathcal{M}(A_1, \dots, A_m)$ gives continuous map of $f : \partial X_n \to G(\mathcal{M}(A_1, \dots, A_m))$. Suppose that f extends to a map $\widehat{f} : X_n \to G(\mathcal{M}(A_1, \dots, A_m))$. Since $\mathcal{M}(A'_1, \dots, A'_m)$ is a retract of $\mathcal{M}(A_1, \dots, A_m)$, we may assume that $\widehat{f}$ maps into $G(\mathcal{M}(A'_1, \dots, A'_m)) = \partial X_n$. But this means that ∂X_n is a retract of X_n, a contradiction. ♦

3.1 Evasive graph properties and Boolean functions

Consider any property $\mathcal{P}_n$ of graphs with n nodes (we assume that if a graph has this property then every graph isomorphic with it also has it). We say that $\mathcal{P}_n$ *trivial*, if either every graph has it or no one has it. A graph property is *monotone* if whenever a graph has it each of its subgraphs has it. For most graph properties that we investigate (connectivity, the existence of a Hamiltonian circuit, the existence of perfect matching, colorability, etc.) either the property itself or its negation is monotone.

We say that a property $\mathcal{P}_n$ of n-node graphs is *evasive*, if every algorithm computing it has to look up, in the worst case, the adjacency of all pairs of nodes. Since we don't care about other resources of the algorithm (time, memory etc.), we can formalize this as follows. We consider a rooted binary tree T where each internal node is labeled by an unordered pair $\{i, j\}$, $1 \leq i, j \leq n$, each edge is labeled by 0 or 1, and each leaf is labeled by 0 or 1. For each internal node, one of its children is connected to it by an edge labeled 1 and the other child, by an edge labeled 0. Given a graph G on $V = \{1, \dots, n\}$, we enter the tree at the root and walk down to a leaf. At an internal node labeled $\{i, j\}$, we we check whether nodes i and j are adjacent in G, and continue on the edge labeled 1 or 0 accordingly. When we reach a leaf, we must see a label 1 if and only it G has property $\mathcal{P}_n$. If this holds for every graph G, we say that T is a decision tree for $\mathcal{P}_n$.

In this language, a property of n-node graphs is *evasive*, if it cannot be computed by a decision tree of depth less than $\binom{n}{2}$. The Aanderraa–Rosenberg–Karp conjecture says that *every nontrivial monotone graph property is evasive.* This conjecture is unsolved, but it has been proved for a number of specific graph properties (connectivity, planarity etc.) Here we prove the following general theorem.

Theorem 3.9 (Kahn, Saks, Sturtevant) *Every nontrivial monotone property of graphs on n nodes, where n is a prime power, is evasive.*

The analogous theorem is also proved for bipartite graphs (Yao). See also `http://www.cs.elte.hu/~lovasz/complexity.pdf` and `http://arxiv.org/abs/cs/0205031` for other special cases, provable by topology or other tools.

Before proving this theorem, it will be useful to generalize it. We consider *Boolean functions* $f:\ \{0,1\}^n \to \{0,1\}$. We say that the function f is *monotone decreasing* (or just "monotone" for this section), if $f(x_1,\ldots,_n) \le f(y_1,\ldots,y_n)$ whenever $x_i \ge y_i$ for all $1 \le i \le n$. A *decision tree* for f is defined just like for a graph property, except that the internal nodes are labeled by the variables. Given a values for the variables, we walk down the tree, turning at each node depending on the value of that variable, and read off the value of the function from the label of the leaf. We say that a Boolean function with n variables is *evasive*, if it cannot be computed by a decision tree of depth less than n.

We define the *automorphism group* of a Boolean function as the set of permutations of the variables that does not change the value of the function. This permutation group is called *transitive*, if for any two variables x and y there is an automorphism that moves x to y.

To see how graph properties fit in, let G be graph with $V = V(G) = \{1,\ldots,n\}$. For every pair $i,j \subseteq V$, let us introduce a Boolean variable x_{ij} with value 1 if i and j are adjacent and 0 if they are not. In this way, any property of n-point graphs can be considered as a Boolean function with $\binom{n}{2}$ variables. This Boolean function has a transitive automorphism group: for any two pairs of nodes there is a permutation of the nodes taking one into the other, which leads to an isomorphic graph and hence does not change the value of the Boolean function.

The following Generalized Aanderaa–Rosenberg–Karp Conjecture is open: *If the automorphism group of a non-constant monotone Boolean function is transitive, then the function evasive.* We prove it in a special case.

Theorem 3.10 *If the automorphism group of a non-constant monotone Boolean function f contains a transitive subgroup Γ, and Γ has a normal subgroup Δ of prime power order such that Γ/Δ is cyclic, then f is evasive.*

Corollary 3.11 *If the automorphism group of a non-constant monotone Boolean function is transitive, and it is a cyclic group, then it is evasive.*

Corollary 3.12 *If the automorphism group of a non-constant monotone Boolean function is transitive, and the number of its variables is a prime power, then it is evasive.*

Next, we describe how evasiveness relates to topology. Let $f:\ \{0,1\}^n \to \{0,1\}$ be a monotone decreasing Boolean function that is not identically 0. For $x \in \{0,1\}^n$, let $\mathrm{supp}(x) = \{i \in \{1,\ldots,n\}:\ x_i = 1\}$. We associate with f the following simplicial complex:

$$\mathcal{K}_f = \big\{\mathrm{supp}(x):\ x \in \{0,1\}^n,\ x \ne 0,\ f(x) = 1\big\}.$$

Lemma 3.13 *If f is non-evasive, then $\mathcal{K}_f$ is contractible.*

Proof. If f is non-evasive, then there is decision tree computing f with depth at most $n-1$. This decision tree starts with checking a variable, say x_1. The two branches of the tree compute the Boolean functions $f_0(x_2,\ldots,x_n) = f(0,x_2,\ldots,x_n)$ and $f_1(x_2,\ldots,x_n) = f(1,x_2,\ldots,x_n)$. Since the depth of these branches is at most $n-2$, the functions f_0 and f_1 are non-evasive. By induction, we may assume that $\mathcal{K}_{f_0}$ and $\mathcal{K}_{f_1}$ are contractible.

Let $\mathcal{K}'_{f_1} = \mathcal{K}_{f_1} \cup \{X \cup \{1\} : X \in \mathcal{K}_{f_1}\}$, then all maximal simplices of $\mathcal{K}'_{f_1}$ contain 1, and so $\mathcal{K}'_{f_1}$ is contractible. We have $\mathcal{K}_{f_0} \cup \mathcal{K}'_{f_1} = \mathcal{K}$ and $\mathcal{K}_{f_0} \cap \mathcal{K}'_{f_1} = \mathcal{K}_{f_1}$, and so Lemma 2.12 implies that $\mathcal{K}$ is contractible. □

Next, we observe the following.

Lemma 3.14 *If f is a nonconstant monotone Boolean function with a transitive automorphism group Γ, then Γ acts on $G(\mathcal{K}_f)$ and has no fixed point.*

Proof. It is clear that every automorphism γ of f is a simplicial map $\gamma : \mathcal{K}_f \to \mathcal{K}_f$, and this extends to a continuous map $\widehat{\gamma} : G(\mathcal{K}_f) \to G(\mathcal{K}_f)$.

We claim that these maps have no common fixed point. Suppose that $x \in G(\mathcal{K}_f)$ satisfies $\widehat{\gamma}(x) = x$ for every $\gamma \in \Gamma$. There is a unique smallest simplex S such that $x \in \mathrm{conv}(S)$. Then $x = \widehat{\gamma}(x) \in \mathrm{conv}(\gamma(S))$, so by the minimality of S, we must have $\gamma(S) = S$. Since the group γ is transitive, this can only happen for every $\gamma \in \Gamma$ if $S = V(\mathcal{K}_f)$. But then the function is identically true. □

Proof of Theorem 3.10. If f is non-evasive, then $G(\mathcal{K}_f)$ is contractible by Lemma 3.13. If f is nontrivial, then Γ acts on $G(\mathcal{K}_f)$ without fixed points. But this contradicts the following theorem of Oliver from topology: *If Γ is a finite group that has a normal subgroup Δ of prime power order such that Γ/Δ is cyclic, then every continuous action of Γ on a contractible simplicial complex has a fixed point.* □

Proof of Theorem 3.9. We consider $V(G)$ as the set of elements of a finite field $\mathbb{F}_q$ (where q is any prime poser), and define Γ as the group of linear transformations of the form $x \mapsto ax+b$, where $a, b \in \mathbb{F}_q$ and $a \neq 0$. Then Γ acts on pairs $\{i, j\} \subseteq \mathbb{F}_q$ transitively. Furthermore, if Δ denotes the subgroup of Γ of transformations of the form $x \mapsto x + b$, then $|\Delta| = q$ is a prime power, and $\Gamma/\Delta \cong F_q^*$ (the multiplicative group of nonzero elements of $\mathbb{F}_q$), which is cyclic by basic results in algebra. So Theorem 3.10 applies and proves the theorem. □

3.2 Topological connectivity

For the next application of Brouwer's Fixed Point Theorem, we need to generalize the notion of contractibility. A topological space is *k-connected* ($k \geq 0$) if for every $0 \leq r \leq k$, every continuous map $f : S^r \to T$ extends to a continuous map $\overline{f} : B^{r+1} \to T$. Equivalently, f is homotopic to a constant map. Sometimes it is convenient to define (-1)-connected as "non-empty". From the results in Section 2.2 it follows that a topological space is contractible if and only if it is k-connected for every k.

An important use of connectivity is the following extension property of maps, which follows along the lines of Lemma 2.8.

Lemma 3.15 *Let $\mathcal{K}$ is a $(k+1)$-dimensional simplicial complex, let $\mathcal{K}'$ be a subcomplex of $\mathcal{K}$, and T a k-connected space. Then every continuous map $f : G(\mathcal{K}') \to T$ extends to a continuous map $G(\mathcal{K}) \to T$.*

Most of the other lemmas from the previous sections can also be generalized, with some care, from contractibility to k-connectivity.

Lemma 3.16 *Let $\mathcal{K}$ be a simplicial complex and $k \geq 0$. Then the following are equivalent:*

(a) *$\mathcal{K}$ is k-connected.*

(b) *The $(k+1)$-dimensional skeleton $\mathcal{K}|_{k+1}$ is k-connected.*

(c) *$\mathcal{K}|_{k+1}$ is a retract of $\Sigma(V(\mathcal{K}))|_{k+1}$.*

Lemma 3.17 *Let $\mathcal{K}$ be a simplicial complex and $U \subseteq V(\mathcal{K})$.*

(a) If $\mathcal{K}$ is k-connected and $\mathcal{K} \cap \Sigma(U)$ is $(k-1)$-connected, then $\mathcal{K} \cup \Sigma(U)$ is k-connected.

(b) If $\mathcal{K} \cup \Sigma(U)$ and $\mathcal{K} \cap \Sigma(U)$ are k-connected, then $\mathcal{K}$ is k-connected.

Lemma 3.18 *Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be two k-connected simplicial complexes and assume that $\mathcal{K}_1 \cap \mathcal{K}_2$ is $(k-1)$-connected. Then $\mathcal{K}_1 \cup \mathcal{K}_2$ is k-connected.*

The Nerve Theorem 2.17 has the following version for k-connectivity.

Theorem 3.19 (Connectivity Nerve Theorem) *Let $\mathcal{K}$ be a simplicial complex and let $\mathcal{K}_1, \dots, \mathcal{K}_m$ be subcomplexes such that $\mathcal{K} = \mathcal{K}_1 \cup \dots \cup \mathcal{K}_m$. Assume that $\mathcal{K}_{i_1} \cap \dots \cap \mathcal{K}_{i_r}$ is either empty or $(k-r+1)$-connected for every $1 \leq r \leq k+1$ and $1 \leq i_1 < \dots < i_r \leq m$. Then $\mathcal{K}$ is k-connected if and only if the nerve of $\{V(\mathcal{K}_1), \dots, V(\mathcal{K}_m)\}$ is k-connected.*

Example 3.20 As an application of this theorem, we analyze the connectivity of the simplicial complex $\mathcal{M}(A_1, \dots, A_m)$ from Example 2.5. If $|A_i| = 1$ for some i, then every maximal simplex contains the single element of A_i, and hence $\mathcal{M}(A_1, \dots, A_m)$ is contractible. If $|A_i| \geq 2$ for every i, then $\mathcal{M}(A_1, \dots, A_m)$ is $(m-2)$-connected. This is trivial if $m \leq 1$, so we may assume that $m \geq 2$. Let $\mathcal{M}_x$ be the set of all simplices in $\mathcal{M}(A_1, \dots, A_m)$ that contain x or can be extended to contain x. Then $\mathcal{M}_x$ is contractible by the argument above. Furthermore,

$$\mathcal{M}(A_1, \dots, A_m) = \bigcup_{x \in A_1} \mathcal{M}_x.$$

Clearly

$$\mathcal{M}_{x_1} \cap \dots \cap \mathcal{M}_{x_r} = \mathcal{M}(A_2, \dots, A_m)$$

if $x_1, \dots, x_r$ are distinct elements of A_1 ($r \geq 2$). The complex $\mathcal{M}_{x_1} \cap \dots \cap \mathcal{M}_{x_r}$ is $(m-3)$-connected by induction. The Connectivity Nerve Theorem 3.19 applies, and gives that $\mathcal{M}(A_1, \dots, A_m)$ is $(m-2)$-connected. ♦

Example 3.21 Let $V = \{1, \dots, n\}$, let $0 \leq a \leq b \leq n$, and consider the bipartite graph G between color classes $U = \binom{V}{a}$ and $W = \binom{V}{b}$, where $X \in U$ is adjacent to $Y \in W$ if and only if $X \subseteq Y$. We claim that the neighborhood complexes $\mathcal{N}_U = \mathcal{N}_U(G)$ and $\mathcal{N}_W = \mathcal{N}_W(G)$ are $(b-a-1)$-connected.

We know that $\mathcal{N}_U \sim \mathcal{N}_W$ (Lemma 2.15), so it suffices to prove that $\mathcal{N}_U$ is $(b-a-1)$-connected. This is trivial in some cases: if $a = b$ (recall that (-1)-connected means non-empty); if $a = 0$ (a one-point space is contractible); and if $b = n$ (a full simplex is contractible). So we may assume that $0 < a < b < n$.

For every $i \in V$, let $U_i = \binom{V\setminus\{i\}}{a}$ and let $\mathcal{K}_i = \mathcal{N}_U \cap \Sigma(U_i)$. Since $b < n$, we have $\mathcal{N}_U = \cup_i \mathcal{K}_i$. By induction on n, we know that $\mathcal{K}_i$ is $(b-a-1)$-connected.

More generally, let $1 \le i_1 < \cdots < i_r \le n$. If $r \le n-b$, then $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r}$ is just the neighborhood complex of the bipartite graph between levels $\binom{V'}{a}$ and $\binom{V'}{b}$, where $V' = V \setminus \{i_1, \dots, i_r\}$. Hence by induction, $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r}$ is $(b-a-1)$-connected. If $n-b \le r \le n-a$, then $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r}$ is a full simplex, so it is contractible. If $r > n-a$, then $\mathcal{K}_{i_1} \cap \cdots \cap \mathcal{K}_{i_r} = \emptyset$. So the Connectivity Nerve Theorem 3.19 applies, and we get that $\mathcal{N}_U$ is $(b-a)$-connected. ♦

Example 3.22 Let P be the poset consisting of levels $a, a+1, \dots, b$ of a finite Boolean algebra on a set V. Adding $\emptyset$ and V, we get a lattice, hence Corollary 2.21 implies that $\mathcal{C}(P)$ is homotopy equivalent to the complex $\mathcal{K}_U$ considered in Example 3.21. In particular, it is $(b-a-1)$-connected. ♦

3.3 Partitioning a graph into connected pieces

Theorem 3.23 *Let G be a k-connected graph on n nodes, $v_1, \dots, v_k \in V$, and $n_1, \dots, n_k$, positive integers with $n_1 + \cdots + n_k = n$. Then there exists a partition $V(G) = V_1 \cup \cdots \cup V_k$ such that $v_i \in V_i$ and $|V_i| = n_i$ for all $1 \le i \le k$.*

It will be convenient to rephrase this theorem as follows. Let G be a simple graph, and let $v \in V(G)$ have degree k. Let $N(v) = \{u_1, \dots, u_k\}$. For every spanning tree T of G and every $1 \le i \le k$, let $n_i(T)$ denote the number of nodes $w \in V(G)$ for which the w-v path in T passes through the edge vu_i (this may be 0 if $vu_i \notin E(T)$). Let $\mathbf{n}(T) = (n_1(T), \dots, n_k(T))$.

Theorem 3.24 *Let G be a k-connected simple graph on n nodes, and let $v \in V(G)$ with $N(v) = \{u_1, \dots, u_k\}$. Let $n_1, \dots, n_k$ be nonnegative integers with $n_1 + \cdots + n_k = n-1$. Then G has a spanning tree T such that $\mathbf{n}_i(T) = n_i$ for all $1 \le i \le k$.*

Let G be a connected (not necessarily simple) graph and let $v \in V(G)$ be a specified "root". Let $S = S(G)$ denote the set of spanning trees of G. We define a simplicial complex $\mathcal{T} = \mathcal{T}(G, v)$ on vertex set S as follows: for a nonempty set $X = \{T_1, \dots, T_r\}$ of spanning trees, let its *kernel* T_X be the connected component of $T_1 \cap \cdots \cap T_r$ containing v (i.e., the largest common subtree of them containing v). Let the set X form a simplex, if every node in $V(G) \setminus V(T_X)$ has degree 1 in every T_i. Note that this implies that every $x \in V(G) \setminus V(T_X)$ is adjacent to a node of T_X in T_i.

Lemma 3.25 *Let G be a k-connected graph and $v \in V(G)$, then $\mathcal{T}(G, v)$ is $(k-2)$-connected for every $v \in V(G)$.*

We need to strengthen the lemma to facilitate induction. First, we allow multiple edges (these would be irrelevant in Theorem 3.24). Second, we need the following relaxation of the condition of k-connectivity:

($\mathbf{A}_k$): every cutset of fewer than k nodes contains v.

Lemma 3.26 *If a multigraph G and $v \in V(G)$ satisfy $(\mathbf{A}_k)$, then $\mathcal{T}(G,v)$ is $(k-2)$-connected for every $v \in V(G)$.*

In particular, if G is k-connected (as a graph), then $\mathcal{T}(G,v)$ is $(k-2)$-connected (as a simplicial complex). As a special case, if v is adjacent to every node of $V(G)\setminus\{v\}$, then $(\mathbf{A}_k)$ is satisfied for every k, and hence $\mathcal{T}(G,v)$ is contractible.

Proof. Let $N(v) = \{u_1, \dots, u_m\}$. Let S_e be the set of spanning trees containing $e \in \nabla(v)$, and set $\mathcal{T}_e = \mathcal{T} \cap \Sigma(S_e)$.

We have to distinguish two cases, depending on whether or not $N(v) = V(G) \setminus \{v\}$. The proof in these two cases will be similar but slightly different.

Case 1. $N(v) = V(G) \setminus \{v\}$. For $S \subseteq \nabla(v)$, let G/S denote the graph obtained from G by contracting every edge e_i to a single node (keeping multiple edges, but we may delete the resulting loops if we wish). We denote by $\overline{v}$ the image of v in such a contraction.

Let U denote the set of spanning trees in which every node except v is a leaf. (Since we allow multiple edges, there may be more than one such tree.) By hypothesis, U is non-empty, and $U \in \mathcal{T}$. Let $\mathcal{T}_0 = \mathcal{T} \cap \Sigma(U)$. We claim that

$$\mathcal{T} = \bigcup_{e \in \nabla(v)} \mathcal{T}_e \cup \mathcal{T}_0. \tag{2}$$

Indeed, if $X \in \mathcal{T}$, then either its kernel T_X has no edge, and then $X \in \Sigma(U)$, or $e \in E(T_X)$ for some $e \in \nabla(v)$, and then $X \in \Sigma(S_e)$.

We verify that the decomposition (2) satisfies the conditions of the Nerve Theorem 2.17. For any set of edges $e_1, \dots, e_r \in \nabla(v)$, $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \neq \emptyset$ if and only if $e_1, \dots, e_r$ go to different nodes in $N(v)$. If this happens, then

$$\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \cong \mathcal{T}(G/\{e_1, \dots, e_r\}, \overline{v}).$$

Since $\overline{v}$ is connected to every node in $G/\{e_1, \dots, e_r\}$, we may assume by induction that $\mathcal{T}(G/\{e_1, \dots, e_r\}, \overline{v})$ is contractible. Hence $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r}$ is contractible. Furthermore, $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \cap \mathcal{T}_0$ is a simplex, and hence contractible.

So the Nerve Theorem 2.17 applies, and we get that $\mathcal{T} \sim \mathsf{nerve}(\{\mathcal{T}_e :\ e \in \nabla(v)\} \cup \{\mathcal{T}_0\})$. But observe that if $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \neq \emptyset$ then also $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \cap \mathcal{T}_0 \neq \emptyset$, which means that $\mathcal{T}_0$ is a vertex of every maximal simplex of this nerve. This implies that the nerve of $\{\mathcal{T}_e :\ e \in \nabla(v)\} \cup \{\mathcal{T}_0\}$ is contractible.

Case 2. $N(v) \neq V(G) \setminus \{v\}$. Then $N(v)$ is a cutset, and condition $(\mathbf{A}_k)$ implies that $m \geq k$. In this case $\{v\}$ cannot be the kernel of any simplex in $\mathcal{T}$, and so

$$\mathcal{T} = \bigcup_{e \in \nabla(v)} \mathcal{T}_e.$$

For $e_1, \dots, e_r \in \nabla(v)$, we have

$$\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \cong \mathcal{T}(G/\{e_1, \dots, e_r\}, \overline{v}).$$

Since $G/\{e_1, \dots, e_r\}$ satisfies condition $(\mathbf{A}_k)$ (with root $\overline{v}$), we may assume by induction that $\mathcal{T}(G/\{e_1, \dots, e_r\}, \overline{v})$ is $(k-2)$-connected. Hence the conditions of the Connectivity Nerve Theorem 3.19 are satisfied, and so it suffices to prove that the nerve of $\{\mathcal{T}_e : \ e \in \nabla(v)\}$ is $(k-2)$-connected.

Observe that $\mathcal{T}_{e_1} \cap \dots \cap \mathcal{T}_{e_r} \neq \emptyset$ if and only if $e_1, \dots, e_r$ connect v to different nodes in $N(v)$. Hence $\mathsf{nerve}\big(\{\mathcal{T}_e : \ e \in \nabla(v)\}\big) \cong \mathcal{M}(E_{u_1 v}, \dots, E_{u_m v})$, where $E_{u_i v}$ is the set of edges connecting v and u_i. This last complex is $(m-2)$-connected, as discussed in Example 3.20. Since $m \geq k$, this completes the proof. □

Proof of Theorem 3.24. We want to prove that if G is k-connected, then for every vector $z \in \mathbb{Z}^k$ such that $z_i \geq 0$ and $\sum_i z_i = n-1$ there is a spanning tree T with $z = \mathbf{n}(T)$.

Let us extend the map $T \mapsto \mathbf{n}(T)$ to a continuous map $g : \ G(\mathcal{T}(G, v)) \to \mathbb{R}^k$ linearly. It suffices to prove the following two facts:

Claim 3.27 *Every point of the simplex $\Delta = \{x \in \mathbb{R}^n : \ x_i \geq 0, \ \sum_i x_i = n-1$ is in the range of g.*

Claim 3.28 *If a point $z \in \Delta \cap \mathbb{Z}^k$ is contained in $f(\mathrm{conv}(S))$ for some simplex $S \in \mathcal{T}(G, v)$, then $z = \mathbf{n}(T_0)$ for some spanning tree T_0.*

Proof of Claim 3.27. For every $\emptyset \neq X \subseteq U = \{1, \dots, k\}$, let $G_X = G \setminus \{u_i v : \ i \notin X\}$. The graph G_X is trivially $|X|$-connected, and hence $S(G_X, v) = \{T \in S(G) : \ n_i(T) = 0 \text{ for } i \notin X\}$ induces a subcomplex $\mathcal{T}(G_X, v) \subseteq \mathcal{T}(G, v)$ that is $(|X|-2)$-connected.

We define a continuous map $f : \ G(\Sigma(U)) \to G(\mathcal{T}(G, v))$ so that for each $S \in \Sigma(U)$, $f(\mathrm{conv}(S)) \subseteq G(\mathcal{T}(G_S, v))$. First, for every vertex $i \in U = V(\Sigma(U))$, let $f(i)$ be any spanning tree T with $n_i(T) = n-1$ (and so $n_j(T) = 0$ for $j \neq i$). Such a tree clearly exists. We extend f step-by-step to the faces of $G(\Sigma(U))$. Let $S \subseteq U$ be a smallest simplex for which f is not yet defined on $\mathrm{conv}(S)$. Then f is defined on $G(\Gamma(S))$, and $f(G(\Gamma(S))) \subseteq G(\mathcal{T}(G_S, v))$. Since $\mathcal{T}(G_S, v)$ is $(|S|-2)$-connected, we can extend f to $\mathrm{conv}(S)$ as desired.

Now the composite map $\mathbf{n} \circ f$ maps the geometric simplex $\mathrm{conv}(U)$ into the geometric simplex Δ so that the image of every face of $\mathrm{conv}(U)$ is contained in the corresponding face of Δ. Corollary 3.5 implies that $\mathbf{n} \circ f$ is surjective, which completes the proof of Claim 3.27.

Proof of Claim 3.28. Let $S = \{T_1, \dots, T_m\}$, and let T be the kernel (largest common subtree) of S. Let W_i be the set of nodes of T separated from the root by the the edge $u_i v$. Let $V(G) \setminus V(T) = \{q_1, \dots, q_t\}$. Let us construct an auxiliary bipartite graph H with color classes $\{u_1, \dots, u_k\}$ and $\{q_1, \dots, q_t\}$, where $u_i q_j \in E(H)$ if and only if there is an edge of G connecting q_j to W_i.

We know that there are real numbers $a_1, \dots, a_m$ such that $a_r \geq 0$, $\sum_r a_r = 1$ and $\sum_r a_r \mathbf{n}(T_r) = z$. Define

$$\beta_{ij} = \sum \{a_r : \ T_r \text{ connects } q_j \text{ to } W_i\}.$$

Then (β_{ij}) is a weighting of the edges of H such that

$$\sum_i \beta_{ij} = 1 \quad (j = 1, \dots, t), \qquad \sum_j \beta_{ij} = z_j - |W_i| \quad (i = 1, \dots, k)$$

Since the matrix of this system of equations is totally unimodular, it has a solution β'_{ij} in nonnegative integers. Connecting q_j to W_i by β'_{ij} edges (this number is trivially 0 or 1), together with T we get a spanning tree T_0 such that $\mathbf{n}(T_0) = z$. $\square$

4 The Borsuk–Ulam Theorem

4.1 Many forms of the Borsuk–Ulam Theorem

Let T be a topological space. An *antipodality* on T is a homeomorphism $\alpha : T \to T$ such that $\alpha \circ \alpha = \mathrm{id}_T$ (α is an involution) and α has no fixed points. Sometimes we denote $\alpha(x)$ by $-x$. The space T, together with the antipodality α, is called an *antipodality space*. An *antipodal map* between antipodality spaces T_1 and T_2 is a continuous map $f : T_1$ to T_2 that satisfies $f(-x) = -f(x)$.

Example 4.1 The unit sphere S^{d-1} in $\mathbb{R}^d$, endowed with the map $x \mapsto -x$, is an antipodality space. The unit ball B^d, however is not: the origin is a fixed point of the map $x \mapsto -x$.

Note that there are many other antipodalities on S^{d-1}. For example, we can project S^{d-1} onto itself from any interior point of the ball. ♦

For a simplicial complex $\mathcal{K}$, an involution $\alpha : V(\mathcal{K}) \to V(\mathcal{K})$ that is simplicial and $\alpha(S) \cap S = \emptyset$ for every $S \in \mathcal{K}$ is called an *antipodality* of $\mathcal{K}$. It is easy to see that every antipodality of $\mathcal{K}$ extends to an antipodality of $G(\mathcal{K})$ that is affine on the faces.

Example 4.2 The convex hull B^d of the set $\{\pm e_i\}$, where $\{e_1, \ldots, e_d\}$ is the standard basis in $\mathbb{R}^d$ is called the *cross-polytope*. The involution $x \mapsto -x$ is an antipodality on the boundary ∂B^d. ♦

Lemma 4.3 *Let $\mathcal{K}$ be a simplicial complex with antipodality α, and let (T, β) be an antipodality space. If $\mathcal{K}$ is d-dimensional and $\mathcal{T}$ is $(d-1)$-connected, then there is an antipodal map $G(K) \to T$.*

Proof. As mentioned before, we can extend α to an antipodality on $G(\mathcal{K})$. We construct the map $f : G(K) \to T$ step by step, starting with the vertices and working up on the dimension. We pick any vertex v of $G(\mathcal{K})$, and let $f(v) \in T$ be an arbitrary point. We are now forced to define then $f(\alpha(v)) = \beta(f(v))$. Continuing, at every step we consider a simplex $S \in \mathcal{K}$ with lowest dimension on which f is not yet defined. Then f is defined on the boundary $\partial G(S)$, and so by the $(d-1)$-connectivity of T, it can be extended to $G(S)$. We are forced now to extend f to $G(\alpha(S))$ by $f(G(\alpha(x)) = \beta(f(x))$. Repeating this, we get the extension to $G(\mathcal{K})$. $\square$

Corollary 4.4 *If T is a k-connected antipodality space, then there exists an antipodal map of $(S^{k+1}, -)$ into T.*

The three assertions in the following theorem are all essentially equivalent to each other, and are called the *Borsuk–Ulam Theorem*.

Theorem 4.5

(a) *There is no antipodal map from* S^r *into* S^{r-1} $(r \geq 1)$.

(b) *For every continuous map* $f: S^r \to \mathbb{R}^r$ *there exists an* $x \in S^r$ *such that* $f(x) = f(-x)$.

(c) *If* S^r *is covered by* $r+1$ *sets, and either each of these is closed or each of these is open, then one of these sets contains an antipodal pair of points.*

4.2 *A polyhedral Borsuk–Ulam Theorem

The Borsuk–Ulam Theorem has the following following discrete version. Let P be a full-dimensional convex polytope in $\mathbb{R}^d$. We say that two faces A and B of P are *opposite* if there exists a linear function $\ell: \mathbb{R}^d \to \mathbb{R}$ that is maximized by the set of points in A and minimized by the set of points in B.

Theorem 4.6 [6] *Let* P *be a full-dimensional convex polytope in* $\mathbb{R}^d$ *and* $f: P \to \mathbb{R}^{d-1}$. *Then* P *has two opposite faces* A *and* B *such that* $f(A) \cap f(B) \neq \emptyset$.

The following slight extension of this theorem will be needed later. We say that the map $f: P \to \mathbb{R}^{d-1}$ is *generic* if for every pair of faces A and B with $\dim(A) + \dim(B) < d-1$, we have $f(A) \cap f(B) = \emptyset$, and for every pair of faces A and B with $\dim(A) + \dim(B) = d-1$, the set $f(A) \cap f(B)$ is finite.

Theorem 4.7 [38] *Let* P *be a full-dimensional convex polytope in* $\mathbb{R}^d$ *and let* $f: P \to \mathbb{R}^{d-1}$ *be a generic map. Then* P *has two opposite faces* A *and* B *such that* $\dim(A) + \dim(B) = d-1$ *and* $|f(A) \cap f(B)|$ *is odd.*

We say that the polytope P in $\mathbb{R}^d$ is *generic*, if for every two opposite faces A and B, $\dim(A) + \dim(B) = d-1$.

Theorem 4.8 *Let* P *be a generic full-dimensional convex polytope in* $\mathbb{R}^d$ *and let* $f: P \to \mathbb{R}^{d-1}$ *be a generic map. Then*

$$\sum |f(A) \cap f(B)| \equiv 1 \pmod{2},$$

where the summation extends over all pairs of opposite faces A *and* B *of* P.

4.3 *A linked Borsuk–Ulam theorem

Let U and W be two disjoint embedded copies of S^r in $\mathbb{R}^{2r+1}$. Then we define their (modulo 2) *linking number* $\ell(U,W)$ as follows. Let $f: S^r \to U$ be a homeomorphism. We extend f to a continuous mapping $F: B^{r+1} \to \mathbb{R}^{2r+1}$ such that the number of points $x \in B^r$ for which $F(x) \in W$ is finite; and then let $\ell(U,W)$ be this number. It can be shown that $\ell(U,W)$ is independent of the choice of f and F, and $\ell(U,W) = \ell(W,U)$.

Theorem 4.9 *Let* P *be a convex polytope in* $\mathbb{R}^{2d+1}$ *and let* f *be an embedding of the* $(d-1)$*-dimensional skeleton of* P *in* $\mathbb{R}^{2d-1}$. *Then there exists a pair of opposite faces* A *and* B *with* $\dim(A) = \dim(B) = d$ *such that* $f(\partial A)$ *and* $f(\partial B)$ *are linked.*

Theorem 4.10 *Let P be a generic convex polytope in $\mathbb{R}^{2d+1}$ and let f be an embedding of the $(d-1)$-dimensional skeleton of P in $\mathbb{R}^{2d-1}$. Then*

$$\sum \ell(\partial A, \partial B) \equiv 1 \qquad (\text{mod}2),$$

where the summation extends over all pairs of opposite faces A and B with $\dim(A) = \dim(B) = d$.

4.4 The Ham Sandwich Theorem

Theorem 4.11 *Let $A_1, \ldots, A_d$ be measurable sets in $\mathbb{R}^d$ with finite measure. Then there exists a closed halfspace H such that $\mu(H \cap A_i) = \frac{1}{2}\mu(A_i)$ for all i.*

Proof. For every vector $v = (v_0, v_1, \ldots, v_d) \in S^d$, we define

$$f_i(v) = \lambda\{x \in A_i : \; v_0 + v_1x_1 + \cdots + v_dx_d \geq 0\} - \lambda\{x \in A_i : \; v_0 + v_1x_1 + \cdots + v_dx_d \leq 0\}$$

(where λ is the Lebesgue measure). The map $v \mapsto (f_1(v), \ldots, f_d(v))$ is antipodal, so by the Borsuk-Ulam Theorem there is a $v \in S^d$ such that $f_1(v) = \cdots = f_d(v) = 0$. But then the halfspace H defined by the inequality $v_0 + v_1x_1 + \cdots + v_dx_d \geq 0$ satisfies the conditions of the Theorem. (We need that the hyperplane $v_0 + v_1x_1 + \cdots + v_dx_d = 0$ has measure 0.) $\square$

4.5 The Necklace problem

Theorem 4.12 *Consider an (open) necklace consisting of pearls of k colors. Suppose that there is an even number of pearls of each color. Then we can split the necklace at k points, and divide the arising pieces between two robbers so that each robber gets exactly half of the pearls of each color.*

Theorem 4.13 (Continuous version) *Let $A_1 \cup \cdots \cup A_k$ be a partition of $[0,1]$ into measurable parts. Then there is a partition $J_1 \cup \cdots \cup J_{k+1}$ of $[0,1]$ into intervals, and a set $I \subseteq \{1, \ldots, k+1\}$ such that*

$$\sum_{r \in I} \lambda(J_r \cap A_i) = \frac{1}{2}\lambda(A_i)$$

for every $i = 1, \ldots, k$.

Proof. For $z \in S^k$, let $J_1(z) \cup \cdots \cup J_{k+1}(z)$ be the partition of $[0,1]$ into intervals with $\lambda(J_i(z)) = z_i^2$. Define

$$f_i(z) = \sum_{r=1}^{k+1} \text{sgn}(z_r)\lambda(A_i \cap J_r(z)) \quad (i = 1, \ldots, k).$$

Then f_i is continuous (this needs a little argument), and $f_i(-z) = -f_i(z)$. Hence by the Borsuk-Ulam Theorem, there is a $z \in S^k$ such that $f_1(z) = \cdots = f_k(z) = 0$. Then $I = \{j : \; \text{sgn}(z_r) = 1\}$ satisfies the requirements of the theorem. $\square$

5 Homomorphisms, chromatic number, and topology

5.1 The Hom set, the Hom graph and the Hom complex

Let F and G be two simple graphs. A map $\varphi: V(F) \to V(G)$ is called a *homomorphism*, if for every edge $ij \in E(F)$ we have $f(i)f(j) \in E(G)$. In this case we write $\varphi: F \to G$. We denote by $\hom(F,G)$ the number of such homomorphisms $F \to G$. This number is a very important parameter in graph theory. We mention two special cases: $\hom(G,K_r)$ is the number of r-colorations of G; $\hom(C_k,G) = \operatorname{tr}(A_G^k) = \sum_i \lambda_i^k$. (Here K_r is the complete r-graph, C_k is the cycle on k nodes, A_G is the adjacency matrix of G, and $\lambda_1, \lambda_2, \dots$ are the eigenvalues of A_G). The first example above shows that to decide whether $\hom(F,G) \neq 0$ is NP-complete.

We go on in a different direction: we define a graph on the set of homomorphisms $F \to G$, by joining $\varphi, \psi: V(F) \to V(G)$ by an edge if $|\{i \in V(F): \varphi(i) \neq \psi(i)\}| = 1$. We denote this graph by $\mathrm{Hom}(F,G)$.

Example 5.1 If G is connected, then $\mathrm{Hom}(K_2,G)$ has at most 2 components. $\mathrm{Hom}(K_2,G)$ is connected if and only if G is not bipartite. ♦

The following theorem (cited without proof), due to Brightwell and Winkler [17], shows that the connectivity of the graphs $\mathrm{Hom}(F,G)$ is connected to the chromatic number of G in the case of higher chromatic numbers as well.

Theorem 5.2 *If* $\mathrm{Hom}(F,G)$ *is connected for every connected graph F with degrees at most $2d$, the $\chi(G) \geq d+2$.*

The set $\mathrm{Hom}(F,G)$ of homomorphisms $F \to G$ can be equipped with a topological structure. We say that a set of homomorphisms $\varphi_1, \dots, \varphi_k: F \to G$ is a *cluster* if for every edge $uv \in E(F)$ and any $1 \leq i < j \leq k$, we have $\varphi_i(u)\varphi_j(v) \in E(G)$. It is clear that these clusters are closed under taking subsets, and hence they form a simplicial complex $\mathcal{H}(F,G)$.

This construction is "functorial", which means that every homomorphism $\psi: G_1 \to G_2$ induces a simplicial map $\widehat{\psi}: \mathcal{H}(F,G_1) \to \mathcal{H}(F,G_2)$ in a canonical way: For every homomorphism $\varphi: F \to G_1$, we define $\widehat{\psi}(\varphi) = \varphi\psi$. It is trivial that this map from $V\big(\mathcal{H}(F,G_1)\big) = \mathrm{Hom}(F,G_1)$ to $V\big(\mathcal{H}\big(F,G_2)\big) = \mathrm{Hom}(F,G_2)$ maps clusters onto clusters. Similarly, $\mathcal{H}(.,.)$ is a "contravariant functor" in its first variable, which means that every homomorphism $\xi: F_1 \to F_2$ induces a simplicial map $\check{\varphi}: \mathcal{H}(F_2,G) \to \mathcal{H}(F_1,G)$.

The complex $\mathcal{H}(K_2,G)$ plays a special role. We will denote it by $\mathcal{H}(G)$. The points of $\mathcal{H}(G)$ can be thought of as ordered pairs (u,v), where u and v are adjacent nodes of G. Each edge of G has two orientations, and so it contributes two vertices to $\mathcal{H}(G)$. A set X of oriented edges forms a simplex if they are the edges of a a complete bipartite subgraph oriented from one bipartition class to the other, or a subset of such a set. The simplicial complex $\mathcal{H}(G)$ has a natural antipodality: $\alpha(u,v) = (v,u)$ defines a simplicial map $\mathcal{H}(G) \to \mathcal{H}(G)$ that extends to a bijective and fixed-point-free involution $\widehat{\alpha}: G(\mathcal{H}(G)) \to G(\mathcal{H}(G))$. If $G_1 \to G_2$ is a graph homomorphism, then the map $\mathcal{H}(K_2,G_1) \to \mathcal{H}(K_2,G_2)$ it induces is antipodal.

Let $G = (V,E)$ be any graph. The *neighborhood complex* $\mathcal{N}(G)$ of G is the simplicial complex consisting of all subsets A of V such that the elements of A have a common neighbor.

For a bipartite graph, the neighborhood complex consists of two components.

Example 5.3 The complex $\mathcal{N}(K_n)$ is homeomorphic to the $(n-2)$-dimensional sphere: $\mathcal{N}(K_n) = \Sigma(V) \setminus \{V\}$ (here $V = V(K_n) = \{1, \ldots, n\}$).

The complex $\mathcal{H}(K_n)$ is homotopy equivalent to the $(n-2)$-dimensional sphere. Define a map $f : G(\mathcal{H}(K_n)) \to \mathbb{R}^n$ as follows. Let $e_1, \ldots, e_n$ be the standard basis in $\mathbb{R}^n$. For $(u,v) \in V(\mathcal{H}(K_n))$, let $f(u,v) = e_v - e_u$. Extend f linearly over $G(\mathcal{H}(K_n))$.

We claim that the origin is not in the range of f. Indeed, consider a point $x \in \mathrm{conv}(f(u_1,v_1), \ldots, f(u_m,v_m))$, where $\{(u_1,v_1), \ldots, (u_m,v_m)\} \in \mathcal{H}(K_n)$. Note that this implies that $u_i \neq v_j$ for $1 \leq i,j \leq m$. If $f(x) = 0$, then there are real numbers $a_i \geq 0$, $\sum_i a_i = 1$, such that $\sum_i a_i f(u_i,v_i) = 0$, and hence

$$\sum_{i=1}^{m} a_i e_{v_i} = \sum_{i=1}^{m} a_i e_{u_i}.$$

Since there are different basis vectors on the left and on the right, this implies that all coefficients are zero, which is a contradiction.

It follows that $x \mapsto f(x)/\|f(x)\|$ is well defined, and gives an antipodal map $G(\mathcal{H}(K_n)) \to S^{n-1}$. Since the range is contained in the hyperplane $H = \{x : \sum_i x_i = 0\}$, this map goes into $S^{n-1} \cap H$, which is a copy of S^{n-2}.

The construction of the inverse map is not given here (we will need this map only). ♦

It is not a coincidence that $\mathcal{H}(K_n)$ and $\mathcal{N}(K_n)$ turned out to be homotopically equivalent, as this is shown by the following lemma.

Lemma 5.4 *For every graph G, the complexes $\mathcal{H}(G)$ and $\mathcal{N}(G)$ are homotopy equivalent.*

Proof. We define a simplicial map $f : V(\mathcal{H}(G)) \to V(\mathcal{N}(G))$ as follows: For every $(u,v) \in V(\mathcal{H}(G))$ we let $f(u,v) = u$. It suffices to show that this map satisfies the conditions of Lemma 2.14: for every simplex $S \in \mathcal{N}(G)$, the complex $\mathcal{H}(G) \cap \Sigma(f^{-1}(S))$ is contractible.

For every $T \in \mathcal{N}(G)$, let $Q(T)$ denote the set of nodes adjacent to every node of T. Then $Q(T) \neq \emptyset$ by definition, and $T \times Q(T) \in \mathcal{H}(G)$. For a given $S \in \mathcal{N}(G)$, we have

$$\mathcal{H}(G) \cap \Sigma(f^{-1}(S)) = \bigcup_{T \in \Sigma(S)} \Sigma(T \times Q(T)), \tag{3}$$

and hence by Corollary 2.16,

$$\mathcal{H}(G) \cap \Sigma(S \times Q(S)) \sim \mathsf{nerve}\{T \times Q(T) : T \in \Sigma(S)\} \tag{4}$$

Now it is easy to see that for $T_1, \ldots, T_r \subseteq S$, we have

$$(T_1 \times Q(T_1)) \cap \cdots \cap (T_r \times Q(T_r)) \neq \emptyset \quad \Leftrightarrow \quad T_1 \cap \cdots \cap T_r \neq \emptyset. \tag{5}$$

Indeed, the implication $\Rightarrow$ is trivial, and if $T_1 \cap \cdots \cap T_r \neq \emptyset$ than $(u,v) \in (T_1 \times Q(T_1)) \cap \cdots \cap (T_r \times Q(T_r))$ for any $u \in T_1 \cap \cdots \cap T_r$ and $v \in Q(S)$. Thus

$$\mathsf{nerve}\{T \times Q(T) : T \in \Sigma(S)\} = \mathsf{nerve}(\Sigma(S)) \sim \Sigma(S).$$

Since $\Sigma(S)$ is contractible, this proves the Lemma. □

Remark 5.5 It would be tempting to try a shorter argument and use the formula

$$\mathcal{H}(G) = \bigcup_{T\in\mathcal{N}(G)} \Sigma(T \times Q(T)),$$

which implies

$$\mathcal{H}(G) \sim \mathsf{nerve}\{T \times Q(T) :\ T \in \mathcal{N}(G)\}.$$

But the nerve on the right is difficult to handle, since (5) does not remain valid if $T_1, \dots, T_r$ are not faces of the same simplex.

5.2 Topological connectivity and chromatic number

Theorem 5.6 [34] *If the neighborhood complex $\mathcal{N}(G)$ of a graph G is k-connected (equivalently, the homomorphism complex $\mathcal{H}(K_2, G)$ is k-connected), then $\chi(G) \ge k + 3$.*

Proof. By Lemma 4.3 there is an antipodal map $S^{k+1} \to \mathcal{H}(K_2, G)$. Suppose that $\chi(G) \le k + 2$, then there is a homomorphism $G \to K_{k+2}$, which induces an antipodal map $\mathcal{H}(K_2, G) \to \mathcal{H}(K_2, K_{k+2})$. As we have seen in Example 5.3, there is an antipodal map $\mathcal{H}(K_2, K_{k+2}) \to S^k$. Composing these maps, we get an antipodal map $S^{k+1} \to S^k$, contradicting the Borsuk-Ulam Theorem. □

The following related theorem (quoted without proof) shows how using the homomorphism complex leads to further results of this type.

Theorem 5.7 [4]. *If $\mathcal{H}(C_{2r+1}, G)$ is k-connected as a topological space for some $r \ge 1$, then the chromatic number of G is at least $k + 4$.*

5.3 Borsuk graphs and Kneser graphs

As an application of the general bound in the previous section, we prove a result about the chromatic number of a a couple of interesting special classes of graphs.

A *Borsuk graph* is defined on a finite subset $V \subseteq S^{d-1}$, where we connect two points by an edge if and only if their spherical distance is at least $2 - 2\varepsilon$ for a given $\varepsilon > 0$ (they are "almost antipodal"). We assume that V is dense enough so that every cap of spherical radius ε contains at least one point of V.

Theorem 5.8 *The chromatic number of the Borsuk graph (as defined above) is at least $d+1$.*

Proof. Let $\alpha :\ V \to \{1, \dots, d\}$ be any coloring of V; we want to prove that there is an edge connecting two nodes with the same color. Let A_i be the set of points of S^{d-1} for which one of closest points in V has color i ($1 \le i \le d$). Then A_i is closed and $\cup_i A_i = S^{d-1}$, so one of the sets A_i (say, A_1) contains two antipodal points x and $-x$ by Theorem 4.5(c). The closest point of V to x is at distance at most ε, and this (or, if there are more such points, one of these) has color 1. Similarly, there is a point of V of color 1 at distance at most ε from $-x$. These two points of color 1 are adjacent. □

Remark 5.9 Borsuk graphs are interesting because they don't contain any odd cycle shorter than $1/\varepsilon$. The next, combinatorially defined graphs have a similar feature.

The *Kneser graph* K_k^n is defined as the graph whose nodes are all k-subsets of an n-set, and two are adjacent if and only if they are disjoint ($n \geq k$).

Theorem 5.10 [34] *For $n \geq 2k$, the chromatic number of the Kneser graph K_k^n is $n-2k+2$.*

Corollary 5.11 *If all k-element subsets of a $(2k+r-1)$-element set are divided into r classes, then one of the classes contains two disjoint k-sets.*

By Theorem 5.6, it suffices to prove the following:

Lemma 5.12 *The $\mathcal{N}(K_k^n)$ is $(n-2k-1)$-connected.*

Proof. The neighborhood complex $\mathcal{N}(K_k^n)$ can be described as follows: its vertices are all k-subsets of $V = \{1,\dots,n\}$, and vertices $S_1,\dots,S_m\}$ form a simplex if any only if $|S_1 \cup \dots \cup S_m| \leq n-k$. It was shown in Example 3.21 that this complex is $(n-2k-1)$-connected. □

Corollary 5.13 *If the k-subsets of an n-set are colored with $n-2k+1$ colors ($n > 2k$), then one of the colors contains two disjoint sets.*

6 *Linklessly embedable graphs

Let $G = (V,E)$ be a graph. An embedding of a graph G in $\mathbb{R}^3$ is called *linkless* if each pair of disjoint circuits in G are unlinked closed curves in $\mathbb{R}^3$. G is *linklessly embedable* if G has a linkless embedding in $\mathbb{R}^3$.

A YΔ transformation of a graph means deleting a vertex of degree 3, and making its three neighbors mutually adjacent. ΔY is the reverse operation (applied to a triangle). The *Petersen family* consists of all graphs obtainable from K_6 by the so-called ΔY- and YΔ-operations. One can check that this family consists of seven graphs (and it includes the Petersen graph).

Theorem 6.1 (Robertson, Seymour and Thomas) *A graph is linklessly embedable if and only if it does not contain any graph in the Petersen family as a minor.*

We call a linear subspace L of $\mathbb{R}^V$ a *connected representation* if for each $x \in L$, $\mathrm{supp}_+(x)$ is nonempty and $\mathrm{supp}_+(x)$ induces a connected subgraph of G. (For a vector $x \in \mathbb{R}^V$, $\mathrm{supp}(x)$ is the *support* of x, that is, $\mathrm{supp}(x) := \{v \in V | x(v) \neq 0\}$. The *positive support* is $\mathrm{supp}_+(x) := \{v \in V | x(v) > 0\}$ and the *negative support* is $\mathrm{supp}_-(x) := \{v \in V | x(v) < 0\}$.)

We define $\lambda(G)$ as the maximum dimension of any connected representation L of G. It is easy to see that $\lambda(G)$ is monotone under taking minors.

Theorem 6.2 (Van der Holst, Laurent and Schrijver) *(a) $\lambda(G) \leq 1$ if and only if G is a forest;*

(b) $\lambda(G) \leq 2$ if and only if G is series-parallel;

(c) $\lambda(G) \leq 3$ if and only if G is a subgraph of a clique sum of planar graphs.

Theorem 6.3 *[38] If G is linklessly embedable, then $\lambda(G) \leq 4$.*

Among the graphs in the Petersen family, K_6 has $\lambda = 5$, but all other graphs have $\lambda = 4$, so the theorem above does not provide a characterization of linklessly embedable graphs. The following invariant was introduced by Colin de Verdière.

Let $G = (V, E)$ be an undirected graph, $V = \{1, \ldots, n\}$. Then $\mu(G)$ is the largest corank of any symmetric real-valued $n \times n$ matrix $M = (m_{i,j})$ satisfying the following conditions:

(i) M has exactly one negative eigenvalue, of multiplicity 1;

(ii) for all $i \neq j$, $m_{i,j} < 0$ if i and j are adjacent, and $m_{i,j} = 0$ if i and j are nonadjacent,

(iii) there is no nonzero symmetric $n \times n$ matrix $X = (x_{i,j})$ such that $MX = 0$ and such that $x_{i,j} = 0$ whenever $i = j$ or $m_{i,j} \neq 0$.

There is no condition on the diagonal entries $m_{i,i}$. (The *corank* $\mathrm{corank}(M)$ of a matrix M is the dimension of its kernel.) Condition (iii) is called the *Strong Arnold Property.*

Theorem 6.4 *[18, 19] (a) $\mu(G) \leq 1$ if and only if G is a path;*

(b) $\mu(G) \leq 2$ if and only if G is outerplanar;

(c) $\lambda(G) \leq 3$ if and only if G is planar.

The following lemma shows that for a matrix M in the definition of $\mu(G)$, $\mathrm{Ker}(M)$ is *almost* a connected representation for G; that this is not always true is shown by the Petersen graph.

For any graph $G = (V, E)$ and $U \subseteq V$, let $N(U)$ be the set of vertices in $V \setminus U$ that are adjacent to at least one vertex in U. For any $V \times V$ matrix and $I, J \subseteq V$, let $M_{I \times J}$ denote the submatrix induced by the rows in I and columns in J, and let $M_I := M_{I \times I}$. For any vector $z \in \mathbb{R}^I$ and $J \subseteq I$, let z_J be the subvector of z induced by the indices in J.

Lemma 6.5 *[25] Let G be a connected graph, let M be a matrix satisfying (i)–(iii), and let x be a vector in $\mathrm{Ker}(M)$ with $G|\mathrm{supp}_+(x)$ disconnected. Then there are no edges connecting $\mathrm{supp}_+(x)$ and $\mathrm{supp}_-(x)$, and each component K of $G|\mathrm{supp}(x)$ satisfies $N(K) = N(\mathrm{supp}(x))$.*

Theorem 6.6 *[38] $\mu(G) \leq 4$ if and only if G is linklessly embedable.*

7 *Equivariant maps

Let G be a finite group acting on a topological space T. We say that the action is *free* if no element of G other than the identity has a fixed point. A *G-space* is a topological spaces with a free G action. We are mostly concerned with the case when $G = \mathbb{Z}_k$, the cyclic group with k elements. For $k = 2$, this specializes to the notion of antipodality spaces.

Example 7.1 Let S^0_k consist of k points in a fixed cyclic order. Then this is a $\mathbb{Z}_k$-space. ♦

Example 7.2 Let S^1_k be S_1, on which $\mathbb{Z}_k$ acts by rotation by $2\pi/k$. ♦

Example 7.3 Let $T^k = \mathrm{conv}\{e_1, \ldots, e_k\}$ and P^{k-1}_k be the boundary of T^k, with $\mathbb{Z}_k$ acting through a cyclic permutation of the coordinates. This action is free if and only if k is a prime. ♦

The *join* $T_1 \vee T_2$ of two topological spaces T_1 and T_2 is obtained by taking $T_1 \times T_2 \times [0,1]$, and shrinking $T_1 \times \{x_2\} \times \{0\}$ to a single point for all $x_2 \in T_2$, and also shrinking $\{x_1\} \times T_2 \times \{0\}$ to a single point for all $x_1 \in T_1$.

Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be two simplicial complexes; assume that $V(\mathcal{K}_1) \cap V(\mathcal{K}_2) = \emptyset$. Then their *join* $\mathcal{K}_1 \vee \mathcal{K}_2$ is defined as the simplicial complex consisting of all sets $A_1 \cup A_2$, $A_i \in \mathcal{K}_i$. It is straightforward that $G(\mathcal{K}_1 \vee \mathcal{K}_2) \approx G(\mathcal{K}_1) \vee G(\mathcal{K}_2)$. The join of boundary complexes of simplicial convex polytopes is the boundary of a simplicial convex polytope.

Free Z_k-actions on two simplicial complexes give free $\mathbb{Z}_k$ action on the join.

Example 7.4 If n is odd, then S_k^n is defined as the join of $(n+1)/2$ copies of S_k^1. If n is even, then S_k^n is the join of $n/2$ copies of S_k^1 and one copy of S_0^k. ♦

Lemma 7.5 *For odd n, S_k^n is homeomorphic with S^n. For even n, S_k^n is homeomorphic with the space obtained by gluing together p copies of B^n along their boundaries.*

Corollary 7.6 *S_k^n is n-dimensional and $(n-1)$-connected.*

Lemma 7.7 *If T is any $(n-1)$-connected space with a $\mathbb{Z}_k$ action, then S_k^n has a covariant map into T.*

Lemma 7.8 *If $\mathcal{K}$ is any n-dimensional simplicial complex with a free $\mathbb{Z}_k$ action, then $G(\mathcal{K})$ has a covariant map into S_k^n.*

Lemma 7.9 *Let n be even. Then every covariant map of S_k^n into itself has index $1 \pmod p$.*

Theorem 7.10 *S_k^n has no covariant map into S_k^{n-1}.*

Corollary 7.11 *If T is any $(n-1)$-connected space with a free $\mathbb{Z}_k$ action, and $\mathcal{K}$ is an $(n-1)$-dimensional simplicial complex with a free $\mathbb{Z}_k$ action, then T has no covariant map into $G(K)$.*

Theorem 7.12 *[8] Let p be a prime. Then for every continuous map $f: S_p^{d(p-1)} \to \mathbb{R}^d$ there is a point $x \in S_p^{d(p-1)}$ such that f is constant on the orbit of x.*

7.1 Kneser hypergraphs

An *r-graph* (or r-uniform hypergraph) is a pair $(V, \mathcal{H})$ where V is a finite set and $\mathcal{H} \subseteq \binom{V}{r}$. The elements of V are called *nodes*, the elements of $\mathcal{H}$ are called *edges*.

A subset $A \subseteq V$ in an r-graph $(V, \mathcal{H})$ is *independent*, if it does not contain any edge. The *chromatic number* of an r-graph $(V, \mathcal{H})$ is the least integer k such that there is a partition $V = A_1 \cup \cdots \cup A_k$ into independent sets.

An *r-box* is an r-graph defined by $V = V_1 \cup \cdots \cup V_r$ and $\mathcal{H} = V_1 \times \cdots \times V_r$, where $V_1, \dots, V_r$ are disjoint finite sets, and the order of elements in edges is ignored.

We define the *box complex* $\mathcal{B}$ of an r-graph $(V, \mathcal{H})$ as the simplicial complex whose vertices are the ordered edges, and where a set of ordered edges forms a simplex if and only if it is a subset of an r-box contained in $(V, \mathcal{H})$. Note that this generalizes the complex $\mathcal{H}(K_2, G)$.

The group $\mathbb{Z}_r$ acts on $\widehat{\mathcal{H}}$ by cyclically permuting the nodes in every ordered edge. So if ω is the generator of $\mathbb{Z}_r$ then $\omega(v_1, \ldots, v_r) = (v_2, \ldots, v_r, v_1)$. Clearly, this action is simplicial on $\mathcal{B}$ and so it defines an action on $G(\mathcal{B})$. It is also easy to see that this action is free.

Theorem 7.13 [3] *If the box complex of an r-graph $(V, \mathcal{H})$ is $t(r-1)-2$-connected, then its chromatic number is at least $t+1$.*

Theorem 7.14 [3] *If all k-element subsets of an $(rk+(r-1)(t-1))$-element set are divided into t classes, then one of the classes contains r disjoint k-sets.*

The k-subsets of a set S with only $rk+(r-1)(t-1)-1$ elements can be divided into t classes so that none of the classes contains r disjoint k-sets. Let $S = S_1 \cup \ldots S_k$, where $|S_1| = \cdots = |S_{t-1}| = r-1$ and $|S_t| = rk-1$. For $i = 1, \ldots, t-1$, let H_i be the set of k-subsets of S intersecting S_i but not $S_1, \ldots, S_{i-1}$. Let H_t be the rest, i.e., the set of all k-subsets of S_t.

The *Kneser hypergraph* $K(n,k,r) = (V, \mathcal{H})$ is defined as follows. Let S be an n-set and $V = \binom{S}{k}$. Let

$$\mathcal{H} = \big\{\{A_1, \ldots, A_r\} : \ A_i \in V,\ A_i \cap A_j = \emptyset \text{ for all } i \neq j\big\}.$$

Then theorem 7.14 (along with the remark following it) can be restated as follows:

Theorem 7.15 *Let $n = rk + (r-1)(t-1)$. Then the chromatic number of $K(n,k,r)$ is $t+1$.*

This is first proved for the case when $r = p$ is a prime. Using theorem 7.13, it suffices to show the following two facts.

Lemma 7.16 *Let p be a prime, and $n = pk + (p-1)(t-1)$. Then the box complex of $K(n,k,p)$ is (at least) $t(p-1)-2$-connected.*

Lemma 7.17 *If Theorem 7.14 is true for two values $r = r'$ and $r = r''$, then it is also true for $r = r'r''$.*

8 Colored Tverberg Theorem

8.1 The chessboard complex

Let $n, m \geq 1$. The *chessboard complex* $\Delta_{m,n}$ is defined on the fields of an $m \times n$ chessboard, and consists of all sets of fields such that no two are in the same row or column.

Lemma 8.1 *[15] The chessboard complex $\Delta_{m,n}$ is (at least) $(\nu-2)$-connected, where*

$$\nu = \min\left\{m, n, \lfloor \frac{m+n+1}{3} \rfloor\right\}.$$

8.2 Colored Tverberg Theorem

Theorem 8.2 *[51] Let p be a prime, $d \geq 1$ and let $A_0, \ldots, A_d \subseteq \mathbb{R}^d$, $|A_i| \geq 2p$. Then one can select from each A_i p distinct points $a_1^i, \ldots, a_p^i \in A_i$ such that*

$$\cap_{j=1}^{p} \mathrm{conv}\{a_j^0, a_j^1, \ldots, a_j^d\} \neq \emptyset.$$

8.3 An application: bisecting hyperplanes

([32, 7, 51])

9 Euler characteristic

The Euler characteristic of a simplicial complex is defined by $\chi(\mathcal{K}) = \sum_{S\in\mathcal{K}}(-1)^{|S|-1}$. This is known to be a topological invariant, i.e., if $\mathcal{K} \cong \mathcal{K}'$, then $\chi(\mathcal{K}) = \chi(\mathcal{K}')$. More generally,

Theorem 9.1 *If $\mathcal{K}_1$ and $\mathcal{K}_2$ are homotopy equivalent, then $\chi(\mathcal{K}_1) = \chi(\mathcal{K}_2)$. In particular, if $\mathcal{K}$ is contractible, then $\chi(\mathcal{K}) = 1$.*

Consider a finite family $\mathcal{F}$ of (non-empty) polytopes (convex, closed, bounded polyhedra) such that

(1) if $Q \in \mathcal{F}$, then every face of Q is in $\mathcal{F}$;

(2) the intersection of any two members of $\mathcal{F}$ is a face of both.

Such a family is called a *convex cell complex*. Clearly the geometric realization of a simplicial complex can be considered as a convex cell complex (where all cells are simplices). All faces of a convex polytope form a convex cell complex.

The union $P = \cup\mathcal{F}$ is called a *polyhedron*, and we call $\mathcal{F}$ a *convex cell decomposition* of P. We define the Euler characteristic of $\mathcal{F}$ by

$$\chi(\mathcal{F}) = \sum_{Q\in\mathcal{F}}(-1)^{\dim(Q)}.$$

It is known that this quantity depends on the union $P = \cup\mathcal{F}$ only, and we denote it by $\chi(P)$.

Example 9.2 The Euler characteristic of any polyhedron homeomorphic to a ball is 1 (in particular, a convex polytope has Euler characteristic 1), of a polyhedron homeomorphic to the d-dimensional sphere is $1 + (-1)^d$. ♦

9.1 The k-equal problem

Given n real numbers $x_1, \dots, x_n$ (called weights), decide if some k of them are equal. Our access to information about the weights is to make pairwise comparisons: is $x_i < x_j$, $x_i = x_j$, or $x_i > x_j$?

This problem has a trivial solution in $O(n \log n)$ steps: just sort the weights, and compare the i-th element with the $(i + k - 1)$-st for every i. For $k = 2$, this is the best we can do: if any sequence of comparisons applied to n distinct weights concludes that they are indeed distinct, then it must sort them, since if the order of two elements is not determined, then these tests allow that they are equal.

For large values of k, the complexity of this problem does decrease. To show this, assume (for simplicity) that $n = 2^m k$. We start with determining the median (the $(2^{m-1}k)$-th largest weight); this takes $O(n)$ comparisons. Then we go on with finding $(2^{m-2}k)$-th largest elements among those smaller and also among those larger than this element (ties are broken arbitrarily). In the j-th phase, those elements found so far split all elements into blocks of size $2^{m-j}k$, and we find the element of each block which splits it into two equal parts (where each element is counted in the block immediately before it).

After m phases, we have found the k-th, $(2k)$-th, ..., $2^m k$-th largest elements. Now if there are k equal elements, then one of these special elements must occur among them; therefore it is enough to compare each of them with $2k$ other elements (in the blocks immediately before and after them) to see if indeed this is the case.

Each phase takes $O(n)$ comparisons, so the total number of comparisons needed is $O(nm) = O(n \log(n/k))$. We are going to prove that (up to a constant factor) this is best possible.

9.2 Linear decision trees

Let P be a polyhedron in $\mathbb{R}^n$. We want to test whether a given vector $x \in \mathbb{R}^n$ belongs to P. A *linear decision tree* is a rooted ternary tree T, where each node v is associated with a linear function $\ell_v(x) = \sum_i a_i x_i + b$, and the three edges connecting an interior node to its descendants are labeled "+", "0" and "−". Starting from the root, we move down the tree; at each internal node v, we check the sign of $\ell_v(x)$ and follow the appropriately labeled edge. Leaves are labeled YES and NO, and arriving at a leaf we read off the answer to the question "is $x \in P$?".

Let $W^-(T)$ and $W^+(T)$ be the sets of NO-leaves and YES-leaves, respectively. Let, for each leaf w, P_w denote the set of inputs leading to leaf w. Each set P_w is a convex subset of $\mathbb{R}^n$, and P is the union of all cells P_w with $w \in W^+$.

To illuminate the connection between linear decision trees and topology, we start with a simple inequality:

Proposition 9.3 *Let $P \subseteq \mathbb{R}^d$ be the union of N disjoint (closed) polyhedra. Then every linear decision tree for P has at least N YES-leaves.*

Proof. Let $P = \cup_{i=1}^N Q_i$, where the Q_i are disjoint polyhedra. For every YES-leaf w, P_w must be contained in one of the Q_i, and hence every Q_i must be the union of one or more polyhedra P_w, where w is a YES-leaf. So there are at least N YES-leaves. □

The following inequality (which is often much sharper) gives a general lower bound on the number of leaves of a linear decision tree for membership in a polyhedron P in terms of its Euler characteristic.

Theorem 9.4 *For every linear decision tree for a bounded polyhedron P in $\mathbb{R}^n$, the number of YES-leaves is at least $|\chi(P)|$.*

Proof. Let T be a linear decision tree for P. Each set P_w is a convex polytope, but not necessarily closed. In fact, P_w is open in its affine hull: the affine hull A_w of P_w is obtained as the intersection of those hyperplanes $l_u(x) = 0$ which tested with equality along the path from the root to w, and the remaining strict inequalities along this path define P_w. We denote by $\bar{P}_w$ the closure of P_w and by ∂P_w, the boundary of P_w in A_w. Clearly $\bar{P}_w$ is a convex polytope and ∂P_w is homeomorphic to sphere.

We have $P = \cup\{P_w : w \in W^+\}$. Unfortunately, the polyhedra $\bar{P}_w$ (even together with their faces) do not form a convex cell decomposition in general. To relate to the Euler characteristic, we consider the following finer decomposition. Our linear decision tree T determines a family of affine hyperplanes $\mathcal{A}_T = \{H_u\}$, where $H_u = \{x \in \mathbb{R}^n : \ell_u(x) = 0\}$ for each inner node u of T. These hyperplanes subdivide $\mathbb{R}^n$ into a number of relatively open convex polyhedra, which we call *cells.* These cells, together with their faces, partition $\mathbb{R}^n$ (points in the same class behave the same way in all tests on the tree). The closures of cells in P form a convex cell decomposition of P, and hence

$$\sum_{C \subseteq P} (-1)^{\dim(C)} = \chi(P).$$

We can partition this sum according to the YES-leaves:

$$\begin{aligned} \chi(P) &= \sum_{w \in W^+} \sum_{C \subseteq P_w} (-1)^{\dim(C)} = \sum_{w \in W^+} \left(\sum_{C \in \bar{P}_w} (-1)^{\dim(C)} - \sum_{C \in \partial P_w} (-1)^{\dim(C)} \right) \\ &= \sum_{w \in W^+} \big(\chi(\bar{P}_w) - \chi(\partial P_w)\big) = \sum_{w \in W^+} (-1)^{\dim(P_w)} \le |W^+|. \end{aligned}$$

Here we have used that the cells contained in $\bar{P}_w$ form a convex cell decomposition of $\bar{P}_w$, he cells of Δ contained in ∂P_w form a convex cell decomposition of ∂P_w, and $\bar{P}_w$ and ∂P_w are homeomorphic to a ball and a sphere, respectively. $\square$

Corollary 9.5 *Every linear decision tree for a bounded polyhedron P in $\mathbb{R}^n$ has depth at least* $\log_3 |\chi(P)|$.

The Euler characteristic of a polyhedron may be small even if its structure is very complicated. For example, if the polyhedron is star-shaped (i.e. it has a point v such that the segment connecting v to any other point is contained in the polyhedron), then its Euler characteristic is 1.

9.3 Mathematical tools

In order to apply our general topological bounds to the k-equal problem, we need some more advanced mathematical tools, which we collect here.

A sequence of polynomials. We sum up some properties of the polynomials

$$p_k(x) = \sum_{j=0}^{k} \frac{x^j}{j!} \tag{6}$$

Let $\alpha_1, \ldots, \alpha_k$ be the roots of p_k, where α_1 has the smallest absolute value. We have $|\alpha_1| < k$, since $\prod_i \alpha_i = \pm k!$. For the derivative we have

$$p'_k(x) = p_{k-1}(x) = p_k(x) - \frac{x^k}{k!},$$

and hence

$$p'_k(\alpha_i) = -\frac{\alpha_i^k}{k!}.$$

It follows that p_k has no multiple roots. We can also use this formula to compute the partial fraction expansion

$$\frac{1}{p_k(x)} = \sum_{i=1}^{k} \frac{1}{p_k'(\alpha_i)(x-\alpha_i)} = k! \sum_{i=1}^{k} \frac{1}{\alpha_i^k (x-\alpha_i)}. \tag{7}$$

We also need the estimate

$$|p_k(x)| \leq \sum_{j=0}^{k} \frac{|x|^j}{j!} \leq \sum_{j=0}^{\infty} \frac{|x|^j}{j!} = e^{|x|}. \tag{8}$$

The Möbius function. Let $(P, \leq)$ be a finite poset, where $|P| = n$. We say that an $n \times n$ real matrix A is a P-matrix, if its rows and columns are indexed by the elements of P, and $A_{xy} = 0$ unless $x \leq y$ in the poset order. It is easy to prove then the sum and product of two P-matrices are P-matrices, and the inverse of an invertible P-matrix is a P-matrix.

We will need three special P-matrices: the identity matrix I, the matrix Z defined by

$$Z_{x,y} = \begin{cases} 1, & \text{if } x \leq y, \\ 0, & \text{otherwise,} \end{cases}$$

and the matrix $M = Z^{-1}$. A usual notation for the entries of M is $M_{x,y} = \mu(x, y)$, where μ is called the *Möbius function* of the poset. It is easy to see that M has integral entries. The basic equations $MZ = I$ and $ZM = I$ can be written as

$$\sum_{x:\ a \leq x \leq b} \mu(a, x) = \begin{cases} 1, & \text{if } a = b, \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

and

$$\sum_{x:\ a \leq x \leq b} \mu(x, b) = \begin{cases} 1, & \text{if } a = b, \\ 0, & \text{otherwise,} \end{cases} \tag{10}$$

for all $a, b \in P$.

Using these conditions, it is easy to verify the formulas given below for the Möbius functions of some basic examples.

Example 9.6 Let L be the lattice of all subsets of a finite set V, then its Möbius function is given by

$$\mu(X, Y) = \begin{cases} (-1)^{|Y \setminus X|}, & \text{if } X \subseteq Y, \\ 0, & \text{otherwise.} \end{cases}$$

♦

Example 9.7 Let L be the lattice of integers $\{1, \ldots, n\}$, in increasing order. Clearly the $x \wedge y = \min\{x, y\}$ and $x \vee y = \max\{x, y\}$. Then the Möbius function is given by

$$\mu(x, y) = \begin{cases} 1, & \text{if } x = y, \\ -1, & \text{if } x = y - 1, \\ 0, & \text{otherwise.} \end{cases}$$

♦

Example 9.8 Let L be the set of integers $\{1, \ldots, n\}$, partially ordered by divisibility. Let $\omega(x)$ denote the number of different prime divisors of the integer x. Then

$$\mu(x, y) = \begin{cases} (-1)^{\omega(y/x)} & \text{if } x \mid y \text{ and } y/x \text{ is square-free,} \\ 0, & \text{otherwise.} \end{cases}$$

We see that $\mu(x, y)$ depends on the ratio y/x only, and so it can be expressed by the classical number-theoretic Möbius function, which is defined as a single-variable function

$$\mu(x) = \begin{cases} (-1)^{\omega(x)} & \text{if } x \text{ is square-free,} \\ 0, & \text{otherwise.} \end{cases}$$

♦

Let $f : P \to \mathbb{R}$ be any function. We define its *summation function* by

$$g(x) = \sum_{y \leq x} f(y).$$

This summation function determines the original:

$$f(x) = \sum_{y \leq x} \mu(0, y) g(y). \tag{11}$$

Expression 11 is called the *Möbius Inversion Formula.* The proof of (11) is immediate if we write these equations and $g = Z^{\top} f$ and $f = M^{\top} g$.

We can express M by Z. Let $U = Z - I$. Every nonzero entry of U is above the diagonal, which implies that $U^n = 0$. We claim that

$$M = \sum_{k=0}^{n-1} (-1)^k U^k. \tag{12}$$

Indeed,

$$Z\Big(\sum_{k=0}^{n-1} (-1)^k U^k\Big) = (I + U)\Big(\sum_{k=0}^{n-1} (-1)^k U^k\Big) = I + (-1)^{n-1} U^n = I.$$

There are many useful identities for the Möbius function, in particular when the poset is a finite lattice L. We only state one (see e.g. [46, 31] for more). Let $a \leq b < 1$, then

$$\sum_{x \wedge b = a} \mu(x, 1) = 0. \tag{13}$$

This can be proved by induction on the number of elements between a and b. If $b = a$, then the identity is a special case of (9). If $a < b$, then we have

$$\sum_{x \geq a} \mu(x, 1) = \sum_{a \leq c \leq b} \sum_{x \wedge b = c} \mu(x, 1).$$

The left side is 0 by (9), while all terms on the right with $c > a$ are 0 by induction. So the term with $c = a$ must be zero as well.

Möbius function and Euler characteristic. For us, the most important consequence of (12) is that the Möbius function is related to the Euler characteristic. Let $(P, \le)$ be a finite poset with Möbius function μ. For $x, y \in P$, let $P_{xy} = \{z \in P : \ x < z < y\}$. Then

$$\mu(x, y) = \chi(\mathcal{C}(P_{xy})) - 1. \tag{14}$$

Indeed, (12) can be written as

$$\mu(x, y) = \sum_{k=0}^{n-1} \sum_{x<z_1<\cdots<z_{k-1}<y} (-1)^k. \tag{15}$$

This shows that the (x, y) entry of U^k counts chains of $k-1$ elements in P_{xy}, so the (x, y) entry on the right side of (12) is just one larger than the Euler characteristic of $\mathcal{C}(P_{xy})$. (The difference of 1 comes from the convention that we did not include the empty set in a simplicial complex.)

In the special case when the poset is a finite lattice L, we can use Corollary 2.21. Let A be the set of atoms of L and $\mathcal{A} = \{X : \ \emptyset \neq X \subseteq A \ : \vee X \neq 1\}$. Then

$$\mu(0, 1) = \chi(\mathcal{C}(L \setminus \{0, 1\})) - 1 = \chi(\mathcal{A}) - 1 = \sum_{X:\ \vee X \neq 1} (-1)^{|X|-1}. \tag{16}$$

Note that we could also write this as

$$\mu(0, 1) = \sum_{X:\ \vee X = 1} (-1)^{|X|}, \tag{17}$$

since trivially

$$\sum_{X:\ \vee X \neq 1} (-1)^{|X|-1} - \sum_{X:\ \vee X = 1} (-1)^{|X|} = \sum_{X} (-1)^{|X|-1} = 0.$$

Lattices of partitions. Let $V = \{1, \dots, n\}$. A *partition* of V is a family $P = \{A_1, \dots, A_k\}$ of disjoint nonempty subsets (partition classes) such that $A_1 \cup \cdots \cup A_k = V$. If $Q = \{B_1, \dots, B_m\}$ is another partition of V, then we say that P is *finer* than Q, in notation $P \le Q$, if every partition class A_i is a subset of one of the classes B_j. We also say that Q is *rougher* than P. The *meet* $P \wedge Q$ of two partitions is the partition consisting of the nonempty intersections $A_i \cap B_j$. The *join* $P \vee Q$ is the finest partition that is rougher than both P and Q (this is uniquely determined). The finest partition of V is the *discreet partition* $\hat{0}$, consisting of singleton sets; the roughest one is the *indiscrete partition* $\hat{1}$, consisting of a single class. To sum up, partitions of V form a lattice, denoted by Π_n.

Let us call a partition *special*, if every partition class is either a singleton or has at least k elements. Let $\Pi_{n,k} \subseteq \Pi_n$ denote the set of special partitions of V. With the same order as in Π_n, the poset $\Pi_{n,k}$ is a lattice. The join of two special partition is the same as their join in Π_n. For $P \in \Pi_n$, let P' denote the partition obtained by splitting every class with fewer than k elements into singletons. Then P' is the unique largest special partition less than P. Based on this, we can describe the meet of two special partitions $P, Q \in \Pi_{n,k}$ as $(P \wedge Q)'$. Note that the two extreme partitions 0 and 1 are elements of $\Pi_{n,k}$.

Generating function of partition numbers. Let $S_k(n,j)$ denote the number of partitions of an n-set into j parts of size at most k. Clearly $S_k(n,j) = 0$ if $n > kj$ or $n < j$. We can choose such a partition by first choosing the size $r \le k$ of a partition class, then choosing the elements in this class, and then partitioning the remaining elements into $j-1$ classes of size at most k. Each partition is counted j times. This gives the recurrence

$$S_k(n,j) = \frac{1}{j}\sum_{r=1}^{k}\binom{n}{r} S_k(n-r, j-1). \tag{18}$$

We need the following formula for the exponential generating function of these numbers:

$$\sum_{n=0}^{\infty} S_k(n,j)\frac{x^n}{n!} = \frac{1}{j!}(p_k(x)-1)^j. \tag{19}$$

This follows from (18) by induction. For $j = 0$ the assertion is trivial. Let $j > 0$, then

$$\begin{aligned}\sum_{n=0}^{\infty} S_k(n,j)\frac{x^n}{n!} &= \frac{1}{j}\sum_{n=0}^{\infty}\sum_{r=1}^{k}\binom{n}{r} S_k(n-r,j-1)\frac{x^n}{n!} \\ &= \frac{1}{j}\sum_{r=1}^{k}\frac{x^r}{r!}\sum_{n=0}^{\infty} S_k(n-r,j-1)\frac{x^{n-r}}{(n-r)!} = \frac{1}{j}p_k(x)\frac{1}{(j-1)!}(p_k(x)-1)^{j-1} \\ &= (p_k(x)-1)^j.\end{aligned}$$

The Möbius function of partition lattices. The Möbius function μ of the partition lattice Π_n is known: for every partition Q,

$$\mu(Q,1) = (-1)^{|Q|-1}(|Q|-1)! \tag{20}$$

This follows by induction. Merging the classes of Q into singletons changes nothing, so we may assume that $Q = 0$. Let B be the partition $\{\{1\},\{2,\dots,n\}\}$. Then

$$\sum_{P:\, B\wedge P=0} \mu(P,1) = 0.$$

It is easy to see that $B \wedge P = 0$ means that either $P = 0$ or P has one class of the form $\{1,i\}$ and the other classes are singletons. In this latter case, $\mu(P,1) = (-1)^{n-2}(n-2)!$ by induction. Thus $\mu(0,1) + (n-1)(-1)^{n-2}(n-2)! = 0$, which proves that $\mu(0,1) = (-1)^{n-1}(n-1)!$.

Using (20), the more general Möbius function values $\mu(P,Q)$ are easy to figure out, but we will not need them.

Our main tool will be a formula for the Möbius function μ_k of $\Pi_{n,k}$. We get there through a sequence of formulas for $\mu_{n,k} = \mu_k(0,1)$ (it would be easy to extend the argument to determine $\mu_k(P,Q)$ in general). Our first expression relates μ_k and μ.

$$\mu_{n,k} = \sum_{\substack{R\in\Pi_n \\ R'=0}} \mu(R,1). \tag{21}$$

Indeed, using the basic equations (9) and (10) for both μ_k and μ,

$$\sum_{\substack{R\in\Pi_n \\ R'=0}} \mu(R,1) = \sum_{R\in\Pi_n} \mu(R,1) \sum_{\substack{Q\in\Pi_{n,k} \\ Q\le R'}} \mu_k(0,Q) = \sum_{Q\in\Pi_{n,k}} \mu_k(0,Q) \sum_{\substack{R\in\Pi_n \\ R\ge Q}} \mu(Q,R) = \mu_k(0,1).$$

(We have used the fact that for a partition R and special partition Q, the relation $Q \leq R'$ is equivalent to $Q \leq R$.) By (20), this can also be written as

$$\mu_{n,k} = \sum_{Q'=0} \mu(Q,1) = \sum_{Q'=0} (-1)^{|Q|-1}(|Q|-1)!. \tag{22}$$

Collecting terms with the same number of classes, we get

$$\mu_{n,k} = \sum_{j=1}^{n} (-1)^{j-1}(j-1)! S_{k-1}(n,j). \tag{23}$$

Our next goal is to determine the generating function

$$F_k(x) = \sum_{n=0}^{\infty} \mu_{n,k} \frac{x^n}{n!}.$$

This can be done using (23) and (19):

$$\begin{aligned} F_k(x) &= \sum_{n=0}^{\infty} \mu_{n,k} \frac{x^n}{n!} = \sum_{n=1}^{\infty} \sum_{j=1}^{n} (-1)^{j-1}(j-1)! S_{k-1}(n,j) \frac{x^n}{n!} \\ &= \sum_{j=1}^{\infty} \sum_{n=j}^{\infty} (-1)^{j-1}(j-1)! S_{k-1}(n,j) \frac{x^n}{n!} = -\sum_{j=1}^{\infty} \frac{1}{j}(1-p_{k-1}(x))^j \\ &= \ln p_{k-1}(x). \end{aligned}$$

From this nice explicit formula for the generating function, we can extract the coefficients. Let $\alpha_1, \dots, \alpha_{k-1}$ be the roots of p_{k-1}, where α_1 has the smallest absolute value. It will be easier to work with the derivative:

$$\begin{aligned} F_k'(x) &= \frac{p_{k-1}'(x)}{p_{k-1}(x)} = 1 - \frac{x^{k-1}}{k! p_{k-1}(x)} = 1 - x^{k-1} \sum_{i=1}^{k-1} \frac{-1}{\alpha_i^{k-1}(x-\alpha_i)} \\ &= 1 - x^{k-1} \sum_{i=1}^{k-1} \frac{1}{\alpha_i^k} \frac{1}{1-(x/\alpha_i)} \\ &= 1 - x^{k-1} \sum_{i=1}^{k-1} \frac{1}{\alpha_i^k} \sum_{h=0}^{\infty} \left(\frac{x}{\alpha_i}\right)^h \\ &= 1 - \sum_{n=k-1}^{\infty} x^{n-1} \sum_{i=1}^{k-1} \alpha_i^{-n}. \end{aligned}$$

Comparing the coefficients of x^n on both sides, we get the following identity, which has been our main goal in this section:

$$\mu_{n,k} = -(n-1)! \sum_{i=1}^{k-1} \alpha_i^{-n}. \tag{24}$$

9.4 Linear decision trees for the k-equal problem

For the k-equal problem, we consider the polyhedron

$$Q = \{x \in [0,1]^n : \text{ there are } k \text{ equal entries in } x\}$$

The algorithm above determines a linear decision tree for Q with depth $O(n \log(n/k))$ and (consequently) with size $(n/k)^{O(n)}$.

Theorem 9.9 *Every linear decision tree for the k-equal-problem has size $(n/k)^{\Omega(n)}$ and (consequently) depth $\Omega(n\log(n/k))$.*

Let $V = \{1,\dots,n\}$. For every set $S \in \binom{V}{k}$, let $L_S = \{x \in [0,1]^n :\ x_{i_1} = \cdots = x_{i_k}\}$; thus $Q = \cup_S L_S$. Clearly L_S is a linear subspace of dimension $n-k+1$, intersected with the unit cube. The k-equal problem is to decide whether $x \in Q$ for points $x \in [0,1]^n$.

We cannot apply Theorem 9.4 directly, because the segment $x_1 = \cdots = x_n$ is contained in every L_S, and so Q is contractible, and so Theorem 9.4 gives a trivial bound. We get around by restricting our interest to weights satisfying $x_1+\cdots+x_{n-1}-(n-1)x_n = 1$ (which excludes the inputs with $x_1 = \cdots = x_n$).

So let $A_S = \{x \in L_S :\ x_1 + \cdots + x_{n-1} - (n-1)x_n = 1\}$, let $\mathcal{H}$ denote the family of all convex sets A_S, and let $P = \cup\mathcal{H}$. If we have a decision tree T to check whether $x \in Q$, then we can check $x \in P$ by checking whether $x_1 + \cdots + x_{n-1} - (n-1)x_n = 1$ holds. This extended decision tree has the same number of YES-leaves. So Theorem 9.4 implies that

Corollary 9.10 *For every decision tree for the k-equal problem, the number of YES-leaves is at least $|\chi(P)|$.*

Our next task is to compute $\chi(P)$.

Theorem 9.11 *The Euler characteristic of P can be expressed in terms of the Möbius function of $\Pi_{n,k}$ as $\chi(P) = \mu_{n,k-1} + 1$.*

Proof. We can apply the Nerve Theorem to the family $\mathcal{H}$, because every intersection of convex sets is either empty or contractible. We get that P is homotopy equivalent to the nerve of the family $\mathcal{H}$.

This nerve is not hard to describe. Every subset $Y \subseteq \mathcal{H}$ can be viewed as a k-uniform hypergraph $H_Y = (V, Y)$.

Claim 9.12 *We have $Y \in \mathcal{N}$ $(Y \subseteq \mathcal{H},\ Y \neq \emptyset)$ if and only if H_Y is disconnected.*

First, suppose that $Y \in \mathsf{nerve}(\mathcal{H})$, then there is a point $x \in \cup_{S\in Y} A_S$. This means that for every $S \in Y$, x_i has the same value. But the equation $x_1 + \cdots + x_{n-1} - (n-1)x_n = 1$ implies that not all x_i have the same value, so H_Y cannot be connected.

Second, suppose that H_Y is disconnected, and let $V = V_1 \cup V_2$ be a partition into nonempty sets such that every set $S \in Y$ is fully contained in one of them. We may assume that $n \in V_2$, and let $r = |V_1|$. The vector

$$x_i = \begin{cases} 1/r & \text{for } i \in V_1, \\ 0, & \text{otherwise.} \end{cases}$$

is then contained in every subspace A_S $(S \in Y)$, and hence $\cap_{S\in Y} A_s \neq \emptyset$, so $Y \in \mathsf{nerve}(\mathcal{H})$. This proves the Claim.

This Claim implies that

$$\chi(P) = \chi(\mathsf{nerve}(\mathcal{H})) = \sum_{\substack{H_Y \text{ disconnected} \\ Y \neq \emptyset}} (-1)^{|Y|-1}. \tag{25}$$

We can translate this expression in terms of the lattice $\Pi_{n,k}$ of special partitions. Let A denote the set of atoms of $\Pi_{n,k}$. Every partition $Q \in A$ consists of a single class of size k and $n-k$ singleton classes; this means that it can be identified with an element of $\binom{V}{k}$. Furthermore, a set $Y \subseteq \binom{V}{k}$ gives rise to a disconnected hypergraph H_Y if and only if the join of the corresponding atoms is not 1. Hence

$$\chi(P) = \sum_{\substack{\emptyset \neq Y \subseteq A \\ \vee Y \neq 1}} (-1)^{|Y|-1} = \mu_{n,k} + 1.$$

proving the theorem. $\square$

Combining with (24), we get

$$\chi(P) = 1 - (n-1)! \sum_{i=1}^{k-1} \alpha_i^{-n}. \tag{26}$$

How large is this expression? The largest term corresponds to the smallest root (in absolute value) of p_k. Thus the largest term has absolute value

$$(n-1)!|\alpha_1|^{-n} > (n-1)!\Big(\frac{1}{k}\Big)^n > \Big(\frac{n}{3k}\Big)^n,$$

which would be good enough for the proof of Theorem 9.9. Unfortunately, it can happen (as shown by numerical computations) that the other, smaller terms in (26) cancel this largest one.

The remedy is to show that this cannot happen for too many consecutive values of n.

Lemma 9.13 *For all n, k with $1 \le k \le n/2$ there exists an integer m such that $n-k+1 \le m \le n$ and $|\mu_{m,k}| > (m-1)!k^{-m-1}$.*

Proof. Consider the polynomial $q(x) = \prod_{i=2}^{k-1}(x-\alpha_i)$. The coefficients can be calculated explicitly from the expansion

$$q(x) = (k-1)!\frac{p_{k-1}(x)}{x-\alpha_1} = -(k-1)!\frac{1}{\alpha_1}\Big(\sum_{t=0}^{k-1}\frac{x^t}{t!}\Big)\Big(\sum_{r=0}^{\infty}\frac{x^r}{\alpha_1^r}\Big)$$

and so for $j < k-1$, the coefficient of x^j is

$$b_j = -(k-1)!\frac{1}{\alpha_1}\sum_{t=0}^{j}\frac{1}{t!}\frac{1}{\alpha_1^{j-t}} = -(k-1)!\alpha_1^{-j-1}p_j(\alpha_1), \tag{27}$$

from which we only need the estimate

$$|b_j| < (k-1)!|\alpha_1|^{-j-1}e^{|\alpha_1|} < k^k|\alpha_1|^{-j-1}.$$

Furthermore, we have

$$\begin{aligned}\sum_{j=0}^{k-2} b_j \frac{\mu_{n-j,k}}{(n-1-j)!} &= -\sum_{j=0}^{k-2}\sum_{i=1}^{k-1} b_j\alpha_i^{-n+j} = -\sum_{i=1}^{k-1} q(\alpha_i)\alpha_i^{-n} \\ &= -q(\alpha_1)\alpha_1^{-n} = -(k-1)!p'_{k-1}(\alpha_1)\alpha_1^{-n} = \alpha_1^{k-1-n}.\end{aligned}$$

Hence there is a j, $0 \le j \le k-1$, such that

$$\left| b_j \frac{\mu_{n-j,k}}{(n-j-1)!} \right| > \frac{1}{k} |\alpha_1|^{k-1-n}.$$

and hence for $m = n - j$,

$$|\mu_{m,k}| > \frac{1}{k-1}(m-1)!|\alpha_1|^{k-1-n}|b_j|^{-1} \ge k^{-k-1}(m-1)!|\alpha_1|^{k-1-m} \ge (m-1)!k^{-m-1}.$$

□

Proof of Theorem 9.9. It is easy to see that for fixed k, the minimum number of YES leaves for the k-equal problem increases with n. For $k \ge n/10$ we need to prove a bound of $\Omega(n)$ for the depth of the tree, which follows by the above remark. So we may assume that $k < n/10$.

Let m be the number in Lemma 9.13, then the number of YES-leaves in a linear decision tree is at least

$$(m-1)!k^{-m-1} - 1 > \left(\frac{m-1}{3k}\right)^{m-1} \ge \left(\frac{n-k}{3k}\right)^{n-k},$$

and hence the depth of the tree is at least

$$(n-k)\log_3 \frac{n-k}{3k} = \Omega\big(n \log \frac{n}{k}\big).$$

References

[1] N. Alon, Splitting necklaces, *Adv. in Math.* **63** (1987), 247–253.

[2] N. Alon, I. Bárány, Z. Füredi, D. Kleitman, *Point selections and weak ϵ-nets for convex hulls*

[3] N. Alon, P. Frankl, L. Lovász, *The chromatic number of Kneser hypergraphs*, Trans. Amer. Math. Soc. **298** (1986), 359–370.

[4] E. Babson and D.N. Kozlov: Proof of the Lovász conjecture, *Annals of Math.* **165** (2007), 965–1007.

[5] R. Bacher, Y. Colin de Verdière, *Multiplicités des valeurs propres et transformations étoile-triangle des graphes*, preprint, 1994.

[6] E.G. Bajmóczy, I. Bárány, On a common generalization of Borsuk's and Radon's theorem, *Acta Mathematica Academiae Scientiarum Hungaricae* 34 (1979) 347–350.

[7] I. Bárány, Z. Füredi, L. Lovász, On the number of halving planes *Combinatorica* **10**, 175–183.

[8] I. Bárány, S.B. Shlosman and A. Szűcs, On a topological generalization of a theorem of Tverberg, *J. London Math. Soc.* **23**, 158–164.

[9] M. Ben-Or: Lower bounds for algebraic computation trees, *15th ACM STOC* (1983) 80–86.

[10] A. Björner, Topological Methods, in: *Handbook of Combinatorics* (ed. R. Graham, M. Grötschel, L. Lovász), Chapter 34, 1819–1872. North-Holland, Amsterdam (1995).

[11] A. Björner, Some Cohen-Macaulay Complexes arising in Group Theory, in *Commutative Algebra and Combinatorics* (ed. M. Nagata and H. Matsumura), Advanced Studies in Pure Mathematics **11**, North-Holland, Amsterdam, 1987, pp. 13–19.

[12] A. Björner, B. Korte, L. Lovász, *Homotopy properties of greedoids*, Advances in Appl. Math. **6** (1985), 447–494.

[13] A. Björner, L. Lovász, A. Yao, Linear decision trees, hyperplane arrangements, and Möbius functions, in: *Proc. 24th ACM SIGACT Symp. on Theory of Computing*, 170–177.

[14] A. Björner, L. Lovász, Linear decision trees, subspace arrangements, and Möbius functions *Journal of the Amer. Math. Soc.* **7**, 677–706.

[15] A. Björner, S. Vreĉica, R. Živaljeviĉ: The connectivity of chessboard complexes, *Journal of the London Math Soc.* **49** (1994), 25–39.

[16] T. Böhme, On spatial representations of graphs, in: *Contemporary Methods in Graph Theory* (R. Bodendieck, ed.), BI-Wiss.-Verl. Mannheim, Wien/Zurich, 1990, pp. 151–167.

[17] G.R. Brightwell, P. Winkler: Graph Homomorphisms and Phase Transitions, *J. Comb. Theory* Series B, **77** (1999), 221–262.

[18] Y. Colin de Verdière, Sur un nouvel invariant des graphes et un critère de planarité, *Journal of Combinatorial Theory, Series B* 50 (1990) 11–21.

[19] Y. Colin de Verdière, On a new graph invariant and a criterion for planarity, in: *Graph Structure Theory* (N. Robertson, P. Seymour, eds.), Contemporary Mathematics, American Mathematical Society, Providence, Rhode Island, 1993, pp. 137–147.

[20] A. Dold: *Lectures on Algebraic Topology*, Springer-Verlag, 1972.

[21] A. Dold: Parametrized Borsuk–Ulam theorems, *Comment. Math. Helv.* **63** (1988) 275–285.

[22] J. Folkman, The homology groups of a lattice, *J. Math. Mech.* **15** (1966) 631–636.

[23] M. Goresky and R. MacPherson: *Stratified Morse Theory*, Ergebnisse, Band 14, Springer-Verlag, Berlin (1988)

[24] E. Győri, On division of graphs to connected subgraphs, in: *Combinatorics I* (eds. A. Hajnal and V.T. Sós), Coll. Math. Soc. J. Bolyai **18** (1987) 485–494.

[25] H. van der Holst, A short proof of the planarity characterization of Colin de Verdière, *Journal of Combinatorial Theory, Series B* 65 (1995) 269–272.

[26] H. van der Holst, M. Laurent, A. Schrijver, On a minor-monotone graph invariant, *Journal of Combinatorial Theory, Series B* 65 (1995) 291–304.

[27] H. van der Holst, L. Lovász, A. Schrijver, On the invariance of Colin de Verdière's graph parameter under clique sums, *Linear Algebra and Its Applications* 226 (1995) 509–517.

[28] J. Kahn, M. Saks and D. Sturtevant, A topological approach to evasiveness, it Combinatorica **4** (1984) 297–306.

[29] D. Kozlov: *Combinatorial Algebraic Topology*, Springer, Berlin, 2008.

[30] M.A. Krasnoselsky and P.P. Zabrejko, *Geometricheskie Zadachi Nelinejnogo Analiza*, Nauka, Moscow, 1975.

[31] L. Lovász: *Combinatorial Problems and Exercises,* Akadémiai Kiadó–North-Holland, Budapest, 1979; Reprinted by AMS Chelsea Publishing, 2007).

[32] L. Lovász, On the number of halving lines, *Annales Univ. Eötvös* **14** (1971), 107-108.

[33] L. Lovász, A homology theory for spanning trees of a graph, *Acta Math. Hung.* **30**, 241-251.

[34] L. Lovász, Kneser's conjecture, chromatic number, and homotopy, *J. Comb. Theory A* **25** (1978), 319-324.

[35] L. Lovász, Topological and algebraic methods in graph theory, in: *Graph Theory and Related Topics*, Academic Press, 1979, 1-14.

[36] L. Lovász, Matroids and Sperner's Lemma, *Europ. J. Combin.* **1**, 65-66.

[37] L. Lovász, Self-dual polytopes and the chromatic number of distance graphs on the sphere, *Acta Sci. Math. Szeged* **45**, 317-323.

[38] L. Lovász, A. Schrijver: A Borsuk theorem for antipodal links and a spectral characterization of linklessly embeddable graphs, *Proceedings of the Amer. Math. Soc.* **126** (1998), 1275–1285.

[39] J. Mather: Invariance of the homology of a lattice, *Proc. Amer. Math. Soc.* **17** (1966) 1120–1124.

[40] J. Matoušek: *Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry*, Springer (2003).

[41] J. Milnor: On the Betti numbers of real algebraic varieties, *Proc. Amer. Math. Soc.* **15** (1964) 275–280.

[42] J.R. Munkres: *Elements of Algebraic Topology*, Addison-Wesley, Menlo Park (1984).

[43] R. Oliver: Fixed-point sets of group actions on finite acyclic complexes, *Comment. Math. Helvet.* **50** (1975), 155–177.

[44] N. Robertson, P.D. Seymour, R. Thomas, A survey of linkless embeddings, in: *Graph Structure Theory* (N. Robertson, P. Seymour, eds.), Contemporary Mathematics, American Mathematical Society, Providence, Rhode Island, 1993, pp. 125–136.

[45] N. Robertson, P. Seymour, R. Thomas, Sachs' linkless embedding conjecture, *Journal of Combinatorial Theory, Series B* 64 (1995) 185–227.

[46] R.P. Stanley: *Enumerative Combinatorics*, Vol. 1, Wadsworth & Brooks/Cole, Monterey (1986).

[47] M. Steele and A. Yao, Lower bounds for algebraic decision trees, *J. Algorithms* **3** (1982), 1–8.

[48] R. Thom: Sur l'homologie des variétés algébraiques réelles, in: *Differential and Algebraic Topology* (ed. S. S. Cairns), Princeton Univ. Press, Princeton (1965).

[49] K. Wagner, Über eine Eigenschaft der ebene Komplexe, *Mathematische Annalen* 114 (1937) 570–590.

[50] R. Živaljević: User's guide to equivariant methods in combinatorics, *Publ Inst. Math. Belgrade* **59(73)** (1996), 114–130.

[51] R.T. Živaljević, S.T. Vrećica: The colored Tverberg's problem and complexes of injective functions, *J. Combinatorial Theory* Ser. A **61** (1991) 309–318.