Topological Methods in Combinatorics and Geometry

Lecture Notes

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

An earlier version of this text was thoroughly revised
by Günter M. Ziegler

Under construction!!!!
Please report errors to: `matousek@kam.mff.cuni.cz`

November 2001

Contents

Preface

There are several combinatorial and geometric results whose proofs (the first proofs and often the only known proofs) involve a surprising application of algebraic topology. Lovász's striking proof of Kneser's conjecture from 1978 was among the first and most prominent examples, dealing with a problem about finite sets with no apparent relation to topology.

During the last two decades, topological methods in combinatorics became more elaborate. On the one hand, quite advanced parts of algebraic topology have been successfully applied. On the other hand, many of the earlier results can now be proved using only fairly elementary topological notions and tools, and while the first topological proofs, like the Lovász' one, are masterpieces of imagination and involve clever problem-specific constructions, reasonably general recipes exist at present. For some types of problems, they suggest how the desired result can be derived from the nonexistence of a certain map ("test map") between two topological spaces (the "configuration space" and the "target space"). Several standard approaches then become available for proving the nonexistence of such a map. Still, the number of different combinatorial results established topologically remains relatively small.

These lecture notes aim at making some of the elementary topological methods more easily accessible to non-specialists in topology. They cover a number of substantial results proved by topological methods, and at the same time they introduce the required material from algebraic topology. Background in undergraduate mathematics is assumed, as well as a certain mathematical maturity, but no prior knowledge of algebraic topology. (But learning more algebraic topology from other sources is certainly encouraged—this text is no substitute for proper foundations of that subject.)

We concentrate on one type of topological tools, namely the Borsuk–Ulam theorem and generalizations. We develop a somewhat systematic theory as far as our very restricted topological means suffice. Other directions, such as applications of Brouwer's fixed point theorem, are not considered here.

History and notes on teaching. These lecture notes started with a course I taught in fall 1993 in Prague; the transcripts of the lectures by the participants served as a basis of the first version, which was published as a technical report (KAM Series 94–272, Charles University, Prague). Some years later, a course partially based on that text was taught by Günter M. Ziegler in Berlin. He made a number of corrections and additions (in the present version, the treatment of Bier spheres in Section 5.5 is based on his writing, and Chapters 1, 2, and 4 bear extensive marks of his improvements). Many discussions with him and his insightful comments have also greatly influenced the present version.

This is a thoroughly rewritten version for a pre-doctoral course I taught in Zürich in fall 2001. Most of the material was covered in the course: Chapter 1 was assigned as an introductory reading text, and the other chapters were presented in approximately 30 hours of teaching (by 45 minutes), with some omissions throughout and only a sketchy presentation of the last chapter.

Sources. The 1994 version of this text was based on research papers, on a thorough survey of topological methods in combinatorics by Björner [Bjö95], and on a survey of combinatorial applications of the Borsuk–Ulam theorem by Bárány [Bár93]. The presentation in the current version benefited greatly from the recent handbook chapter by Živaljević [Živ97] ([Živ96] is an extended version). The continuation [Živ98] of that chapter deals with more advanced methods beyond the scope of this text.

For learning algebraic topology, many textbooks are available (although in this difficult subject it is probably much better to attend good courses). The first steps can be made with Munkres [Mun00] (which includes preparation in general topology) or Stillwell [Sti93]. A very good and reliable basic textbook is Munkres [Mun84], and Hatcher [Hat01] is a vividly written modern book reaching to quite advanced material in some directions.

Acknowledgments. As was already mentioned, a large contribution to this text was made by Günter M. Ziegler. For answers to my numerous questions I am indebted to Rade Živaljević, Imre Bárńy, and Anders Björner. The participants of the courses (in Prague and in Zürich) provided a stimulating teaching environment and many valuable comments. The end-of-proof symbol is based on a photo of the European badger ("borsuk" in Polish) by Steve Jackson, and it used with his kind permission.

Zürich, November 2001 *Jiří Matoušek*

Preliminaries

This section summarizes rather standard mathematical notions and notation and it serves mainly for reference. More special notions are introduced gradually later on.

Sets. If S is a set, $|S|$ denotes the number of elements (cardinality) of S. By 2^S we denote the set of all subsets of S (the powerset), and $\binom{S}{k}$ stands for the set of all subsets of S of cardinality exactly k. We use $[n]$ to denote the finite set $\{1, 2, \ldots, n\}$.

The letters $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Q}$, and $\mathbb{Z}$ stand for the real numbers, the complex numbers, the rational numbers, the integers, respectively.

Geometry. The symbol $\mathbb{R}^d$ denotes the Euclidean space of dimension d. Points in $\mathbb{R}^d$ are typeset in boldface and they are understood as row vectors; thus, we write $\boldsymbol{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$. The *scalar product* of two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d$ is $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \boldsymbol{x}\boldsymbol{y}^T = x_1y_1 + x_2y_2 + \cdots + x_dy_d$. The *Euclidean norm* of $\boldsymbol{x}$ is $\|\boldsymbol{x}\| = \sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle} = \sqrt{x_1^2 + \cdots + x_d^2}$. Occasionally we also encounter the ℓ_p-norm $\|\boldsymbol{x}\|_p = (|x_1|^p + |x_2|^p + \cdots + |x_d|^p)^{1/p}$, $1 \leq p < \infty$, or the ℓ_∞-norm (or maximum norm) $\|\boldsymbol{x}\|_\infty = \max\{|x_1|, |x_2|, \ldots, |x_d|\}$.

A *hyperplane* in $\mathbb{R}^d$ is a $(d-1)$-dimensional affine subspace, i.e. a set of the form $\{\boldsymbol{x} \in \mathbb{R}^d : \langle \boldsymbol{a}, \boldsymbol{x} \rangle = b\}$ for some nonzero $\boldsymbol{a} \in \mathbb{R}^d$ and some $b \in \mathbb{R}$. A (closed) *halfspace* has the form $\{\boldsymbol{x} \in \mathbb{R}^d : \langle \boldsymbol{a}, \boldsymbol{x} \rangle \leq b\}$, with $\boldsymbol{a}$ and b as before.

The unit ball $\{x \in \mathbb{R}^d : \|x\| \leq 1\}$ is denoted by B^d, while $S^{d-1} = \{x \in \mathbb{R}^d : \|x\| = 1\}$ is the $(d-1)$-dimensional unit sphere (note that S^2 lives in $\mathbb{R}^3$!).

A set $C \subseteq \mathbb{R}^d$ is *convex* if for every $\boldsymbol{x}, \boldsymbol{y} \in C$, the segment $\boldsymbol{x}\boldsymbol{y}$ is contained in C. The *convex hull* of a set $X \subseteq \mathbb{R}^d$ is the intersection of all convex sets containing X and it is denoted by $\operatorname{conv}(X)$. Each point $\boldsymbol{x} \in \operatorname{conv}(X)$ can be written as a *convex combination* of points of X: there are points $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_n \in X$ and real numbers $\alpha_1, \ldots, \alpha_n \geq 0$ such that $\sum_{i=1}^n \alpha_i = 1$ and $\boldsymbol{x} = \sum_{i=1}^n \alpha_i \boldsymbol{x}_i$ (if $X \subseteq \mathbb{R}^d$ we can always choose $n \leq d+1$).

A *convex polytope* is the convex hull of a finite point set in $\mathbb{R}^d$. Each convex polytope can also be expressed as the intersection of finitely many halfspaces. Conversely, if an intersection of finitely many halfspaces is bounded, then it is a convex polytope. A *face* of a convex polytope P is either P itself or an intersection $P \cap h$, where h is a hyperplane that does not dissect P (i.e. not both of the open halfspaces defined by h may intersect P).

Graphs and hypergraphs. Graphs are considered simple and undirected unless stated otherwise, so a graph G is a pair (V, E), where V is a set (the *vertex set*) and $E \subseteq \binom{V}{2}$ is the *edge set*. For a given graph G, we write $V(G)$ for the vertex set and $E(G)$ for the edge set. A *complete graph* has all possible edges, i.e. it is of the form $(V, \binom{V}{2})$. A complete graph on n vertices is denoted by K_n. A graph G is *bipartite* if the vertex set can be partitioned into two disjoint subsets V_1 and V_2, the *(color) classes*, so that each edge connects a vertex of V_1 to a vertex of V_2. A *complete bipartite graph* $K_{m,n}$ has $|V_1| = m$, $|V_2| = n$, and $E = \{\{v_1, v_2\} : v_1 \in V_1, v_2 \in V_2\}$ (so $|E| = mn$).

A *hypergraph* is a pair (V, E), where V is a (usually finite) set and $E \subseteq 2^V$ is a system of subsets of V. The elements of E are called the *edges* or *hyperedges*. A hypergraph is the same thing as a set system but calling it a hypergraph emphasizes a "graph-theoretical" point of view; many notions concerning graphs have natural analogues for hypergraphs.

A hypergraph is *k-uniform* if all of its edges have cardinality k. A hypergraph (V, E) is *k-partite* if there is a partition $V = V_1 \dot\cup V_2 \dot\cup \cdots \dot\cup V_k$ such that $|e \cap V_i| \leq 1$ for every $e \in E$ and every $i \in [k]$.

Miscellaneous. The notation $a := B$ means that the expression B defines the symbol a.

For a real number x, $\lfloor x \rfloor$ denotes the largest integer $\leq x$ and $\lceil x \rceil$ means the smallest integer $\geq x$.

1

Simplicial Complexes

Here we introduce elementary concepts of algebraic topology indispensable for the subsequent chapters, most notably geometric and abstract simplicial complexes, homotopy, and homotopic equivalence of spaces.

Simplicial complexes provide a link from combinatorics to topology. Suppose that we investigate some combinatorial object. Whenever we associate a hereditary set system to our object, we have also associated a topological space—the polyhedron of the corresponding simplicial complex. This space can be studied by methods of algebraic topology, and often its topological properties are linked to combinatorial properties of the original object in interesting ways. Of course, creating simplicial complexes at every possible occasion is no panacea, but sometimes it does lead to meaningful results.

Most of the material of this chapter is usually covered in introductory courses of algebraic topology. But our presentation may deviate from others in details of notation and terminology and it also includes some less commonly treated results, and so even experts in algebraic topology may want to go through the chapter quickly.

1.1 Topological spaces

Although this may be unnecessary for most readers, we first review a few concepts from general topology. We begin with recalling the definition of a topological space, which is a mathematical structure capturing the notions of "nearness" and "continuity" on a very general level.

1.1.1 Definition. *A* **topological space** *is a pair* $(X, \mathcal{O})$*, where* X *is a (typically infinite) ground set and* $\mathcal{O} \subseteq 2^X$ *is a set system, whose members are called the* **open sets***, such* $\emptyset \in \mathcal{O}$, $X \in \mathcal{O}$*, the intersection of finitely many open sets is an open set, and so is the union of an arbitrary collection of open sets.*

Every subset $Y \subseteq X$ defines a subspace, namely the topological space $(Y, \{U \cap Y : U \in \mathcal{O}\})$.

If $(X_1, \mathcal{O}_1)$ and $(X_2, \mathcal{O}_2)$ are topological spaces, a mapping $f: X_1 \to X_2$ is called *continuous* if preimages of open sets are open, i.e. $f^{-1}(V) \in \mathcal{O}_1$ for every $V \in \mathcal{O}_2$.

We implicitly assume that *all the considered mappings between topological spaces are continuous*, although we do not always explicitly say so. More precisely, this applies for unspecified mappings in statements like "let $f\colon S^n \to \mathbb{R}^n$ be a mapping...;" sometimes, after having constructed some mapping, we have to verify its continuity.

What spaces are we going to encounter? The theory dealing with topological spaces in general, *point-set topology* or *general topology*, often investigates fairly exotic examples. However, in our text, as well as in most of algebraic topology, one deals only with topological spaces which are subspaces of some $\mathbb{R}^d$, or at least can be identified with such subspaces.

As the reader certainly knows, a set $U \subseteq \mathbb{R}^d$ is open if for every $x \in U$ there is some $\varepsilon > 0$ such that the ball $\{y \in \mathbb{R}^d : \|x - y\| < \varepsilon\}$ is contained in U. Now let $X \subseteq \mathbb{R}^d$ be an arbitrary set. What are the open sets in the topology of the subspace defined by X? They are exactly the intersections of open sets in $\mathbb{R}^d$ with X; note that they need *not* be open as subsets of $\mathbb{R}^d$ (take X as a closed segment in $\mathbb{R}^2$, for example).

Let us remark that if X is a set and the topology on X is understood, say when $X \subseteq \mathbb{R}^d$ and X is considered with the subspace topology, one usually does not mention the topology in the notation and writes "topological space X" even when formally X is only a set. We will also often say just "space" instead of "topological space."

The topology of $\mathbb{R}^d$, as well as of its subspaces, is induced by a metric, namely by the usual Euclidean metric, which for many readers may be a notion more familiar than topology. But in the considerations of algebraic topology, the metric plays only auxiliary role: often it is a convenient tool but ultimately it is only the topology of a space that really matters. Two spaces that look metrically quite different can be topologically the same; an example are the real line $\mathbb{R}$ and the open interval $(0, 1)$.

In the formulation of some topological definitions and theorems, it would be artificial to restrict to subspaces of Euclidean spaces. But everywhere we assume that the considered spaces are (at least) *Hausdorff*, meaning that for every two distinct points $x, y \in X$ there are disjoint open sets U, V with $x \in U$ and $v \in V$.

Homeomorphism. The notion of "being the same" for topological spaces is similar to many other mathematical structures, such as groups, rings, graphs, and so on. For most mathematical structures, one speaks about isomorphism, which is a bijective mapping preserving the considered structure (group or ring operations, graph edges, etc.). For topological spaces, the corresponding notion is traditionally called a *homeomorphism*.

1.1.2 Definition. *A* **homeomorphism** *of topological spaces $(X_1, \mathcal{O}_1)$ and $(X_2, \mathcal{O}_2)$ is a bijection $\varphi\colon X_1 \to X_2$ such that for every $U \subseteq X_1$, $\varphi(U) \in \mathcal{O}_2$ if and only if $U \in \mathcal{O}_1$. In other words, a bijection $\varphi\colon X_1 \to X_2$ is a homeomorphism if and only if both φ and φ^{-1} are continuous.*

(Warning: there are examples of continuous bijections for which the inverse mapping is not continuous, so both conditions need checking in general.)

If X and Y are topological spaces and there is a homeomorphism $X \to Y$, we write $X \cong Y$ (read "X is homeomorphic to Y").

Closure, boundary, interior. A set F in a topological space X is *closed* iff $X \setminus F$ is open. The *closure* of a set $Y \subseteq X$, denoted by $\mathrm{cl}_X Y$, is the intersection of all closed sets in X containing Y (the subscript X is omitted if X is understood). For $Y \subseteq X = \mathbb{R}^d$, we have $\mathrm{cl}\, Y = \{x \in \mathbb{R}^d : \mathrm{dist}(x, Y) = 0\}$, where $\mathrm{dist}(x, Y) = \inf\{\|x - y\| : y \in Y\}$. The *boundary* of Y is $\partial Y = \{\mathrm{cl}\,(Y) \cap \mathrm{cl}\,(X \setminus Y)\}$and the *interior* $\mathrm{int}\, Y = Y \setminus \partial Y$.

Compactness. We conclude this nano-course of general topology by recalling compactness. A space $X \subseteq \mathbb{R}^d$ is compact if and only if X is a closed and bounded set. (In general, a topological space X is compact if for every collection $\mathcal{U}$ of open sets with $\bigcup \mathcal{U} = X$, there exists a finite $\mathcal{U}_0 \subseteq \mathcal{U}$ with $\bigcup \mathcal{U}_0 = X$.) In a compact metric space, any infinite sequence has a convergent subsequence.

If X is a compact space and $f: X \to \mathbb{R}$ is a continuous real function, then f attains its minimum (and maximum); that is, there is an $x \in X$ with $f(x) \leq f(y)$ for all $y \in X$. Moreover, a continuous function on a compact metric space is *uniformly continuous*; that is, for every $\varepsilon > 0$ there is a $\delta > 0$ such that any two points at distance at most δ are mapped to points at distance at most ε.

Notes. Among many textbooks of topology, we mention Munkres [Mun00] which deals both with general topology and with elements of algebraic topology. A large menagerie of topological spaces is collected in [SS78].

Exercises

1. Verify the following homeomorphisms:

   (a) $\mathbb{R} \cong (0, 1) \cong (S^1 \setminus \{(0, 1)\})$;

   (b) $S^1 \cong \partial([0, 1]^2)$.

2. (a) Let X and Y be topological spaces. Check that a mapping $f: X \to Y$ is continuous if and only if $f^{-1}(F)$ is closed for every closed set $F \subseteq Y$.

   (b) Let X be covered by finitely many closed sets $A_1, A_2, \ldots, A_n$ (i.e. $X = A_1 \cup A_2 \cup \cdots \cup A_n$), and let $f: X \to Y$ be a mapping whose restriction to each A_i is continuous. Verify that f is continuous.

1.2 Homotopy equivalence and homotopy

In algebraic topology, two spaces are considered "the same" under an equivalence relation even coarser than homeomorphism. This notion is called homotopy equivalence. Similarly, continuous maps are classified into classes according to so-called homotopy.

Before plunging into subtleties of homotopy equivalence, we introduce the perhaps more intuitive notion of deformation retract. The horizontal figure 8 drawn by thick line is a deformation retract of the gray area with two holes:

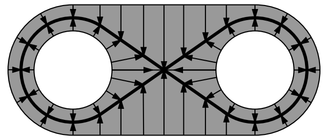

This means that the gray area can be continuously shrunk to the figure 8 while keeping the points of the 8 fixed. The motion is indicated by arrows: each point moves in the shown direction at uniform speed until it hits the 8, where it stops. In general, if X is a space and $Y \subseteq X$ a subspace of it, a *deformation retraction* of X onto Y is a family $\{f_t\}_{t\in[0,1]}$ of continuous maps $f_t: X \to X$ (we can think of t as time), such that f_0 is the identity map on X, $f_t(y) = y$ for all $y \in Y$ and all $t \in [0,1]$ (Y remains stationary), and $f_1(X) = Y$. Moreover, the mappings should depend continuously on t. That is, if we define the mapping $F: X \times [0,1]$ by $F(x,t) = f_t(x)$, this mapping should be continuous. Explicitly, this means that if we choose $x \in X$, $t \in [0,1]$, and an arbitrarily small neighborhood V of $F(x,t)$, there are $\delta > 0$ and a neighborhood U of x such that $F(x', t') \in V$ for all $x' \in U$ and all $t' \in (t+\delta, t-\delta)$. In most of the literature, a deformation retraction is formally viewed as the mapping F, rather than a family of maps; we will use both these presentations interchangeably.

If a deformation retraction exists, Y is called a *deformation retract* of X.

If Y is a deformation retract of X, then X and Y are homotopy equivalent. But, obviously, being a deformation retract is not an equivalence relation; for example, the three black figures below are all deformation retracts of the same gray area as above, but it can be proved that none of them is a deformation retract of another:

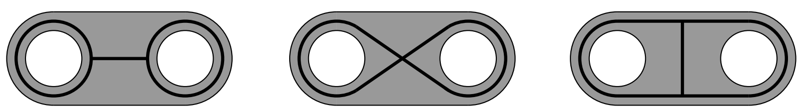

Homotopy equivalence can be introduced as follows: spaces X and Y are homotopy equivalent, in symbols $X \simeq Y$, iff there exists a space Z such that both X and Y are deformation retracts of Z.

The usual definition of homotopy equivalence is different; it is technically more convenient but perhaps less intuitive. To state it, we first need to introduce homotopy of maps.

1.2.1 Definition. *Two continuous maps $f, g: X \to Y$ are* **homotopic** *(written $f \sim g$) if there is a "continuous interpolation" between them; that is, a family $\{f_t\}_{t\in[0,1]}$ of maps $f_t: X \to Y$ depending continuously on t (i.e. the associated bivariate mapping $F(x,t) := f_t(x)$ is a continuous map $X \times [0,1] \to Y$, similar to deformation retraction above) such that $f_0 = f$ and $f_1 = g$.*

In particular, a map $X \to Y$ is called *nullhomotopic* if it is homotopic to a constant map that maps all of X to a single point $y_0 \in Y$ (so "nullhomotopic" is a misnomer; it would be more logical to say "constant-homotopic," but we stick to the traditional terminology). It is not hard to verify that "being homotopic" is an equivalence on the set of all continuous maps $X \to Y$.

1.2.2 Definition (Homotopy equivalence). *Two spaces X and Y are homotopy equivalent (or have the same* **homotopy type***) if there are continuous maps $f: X \to Y$ and $g: Y \to X$ such that the composition $f \circ g: Y \to Y$ is homotopic to the identity map id_Y and $g \circ f \sim \mathrm{id}_X$.*

The equivalence of this definition to the characterization above (homotopy equivalent spaces are deformation retracts of the same space) is nontrivial; see e.g. [Hat01, Chapter 0].

A space homotopy equivalent to a single point is called *contractible*. Some spaces are "obviously" contractible, such as the ball B^d, but for others, contractibility is not easy to visualize. A beautiful example of this is "Bing's house;" see [Hat01, Chapter 0] for a nice presentation. It is tempting to think that a contractible space can always be deformation-retracted to a point, but this is false in general (it can happen that all points are forced to move during any contraction; see Exercise 6).

The task of determining whether two given spaces are homotopy equivalent or not is in general very difficult. Without a sophisticated technical apparatus, it is quite hard to prove even "obvious" facts such as that the circle S^1 is not contractible. But the spaces arising in many topological proofs of combinatorial or geometric theorems happen to be relatively simple, and often they turn out to be homotopy equivalent to a sphere.

Exercises

1. Show that the dumbbell ○-○ and the letter θ are homotopy equivalent, using Definition 1.2.2 (exhibit suitable mappings f and g).

2. Take a 2-dimensional sphere (in $\mathbb{R}^3$) and connect the north and south poles by a segment, obtaining a space X. Let Y be a 2-dimensional sphere with a circle attached by one point to the north pole of the sphere. Show that $X \simeq Y$ (using both the definitions of homotopy equivalence given in the text).

3. Consider two embeddings f and g of the circle S^1 into $\mathbb{R}^3$, where f just inserts the circle into $\mathbb{R}^3$ without changing its shape while g maps it to the trefoil knot

   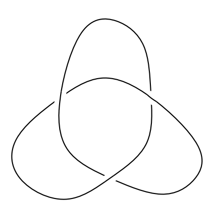

   Are f and g homotopic or not? Substantiate your answer at least informally.

4. (a) Prove that homotopy is an equivalence on the set of all continuous maps $X \to Y$.

(b) Prove that homotopy equivalence is indeed an equivalence on the class of all topological spaces (check transitivity).

5. (a) Prove that a space X is contractible if and only if for every space Y and every continuous map $f: X \to Y$, f is nullhomotopic.

   (b) Prove that a space X is contractible if and only if for every space Y and every continuous map $f: Y \to X$, f is nullhomotopic.

6. The *topologist's comb* is the subspace $X := (R \times [0,1]) \cup ([0,1] \times \{0\})$ of $\mathbb{R}^2$, where R denotes the set of all rational numbers in the interval $[0,1]$. Let Y be made of countably many copies of X arranged in a zigzag fashion into a doubly infinite chain:

   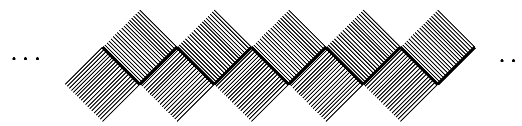

   Show that Y is contractible.

   It can be proved that no point is a deformation retract of Y (you may want to try this as well). In $\mathbb{R}^3$, one can even construct a contractible compact Y with this property; see the exercises to Chapter 0 in Hatcher [Hat01].

1.3 Geometric simplicial complexes

Many topologically interesting subspaces of $\mathbb{R}^d$ can be described as simplicial complexes. This means that they are pasted together from simple building blocks, called simplices and including segments, triangles, and tetrahedra, in a way respecting simple rules. As we will see later, simplicial complexes have a purely combinatorial description and they are particularly significant in the interplay of topology and combinatorics.

First we need to introduce affine independence and simplices.

1.3.1 Definition. *Let $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ be points in $\mathbb{R}^d$. We call them* **affinely independent** *if there are no real numbers $\alpha_0, \alpha_1, \ldots, \alpha_k$, not all of them 0, such that $\sum_{i=0}^k \alpha_i \boldsymbol{v}_i = \mathbf{0}$ and $\sum_{i=0}^k \alpha_i = 0$.*

For $k = 2$, affine independence simply means $\boldsymbol{v}_0 \neq \boldsymbol{v}_1$, for $k = 3$ it means that $\boldsymbol{v}_0, \boldsymbol{v}_1, \boldsymbol{v}_2$ do not lie on a common line, for $k = 4$ it means that $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_3$ do not lie on a common plane, and so on.

Here are two further, simple but useful characterizations of affine independence.

1.3.2 Lemma. *Both of the following conditions are equivalent to affine independence of points $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \mathbb{R}^d$:*

- *The k vectors $\boldsymbol{v}_1 - \boldsymbol{v}_0, \boldsymbol{v}_2 - \boldsymbol{v}_0, \ldots, \boldsymbol{v}_k - \boldsymbol{v}_0$ are linearly independent.*

- *The $(d+1)$-dimensional vectors $(1, \boldsymbol{v}_0), (1, \boldsymbol{v}_1), \ldots, (1, \boldsymbol{v}_k) \in \mathbb{R}^{d+1}$ are linearly independent.*

We leave the easy proof as a warm-up exercise. Also note that $d+1$ is the largest size of an affinely independent set of points in $\mathbb{R}^d$.

Simplices. Here are examples of simplices: a point, a line segment, a triangle, and a tetrahedron:

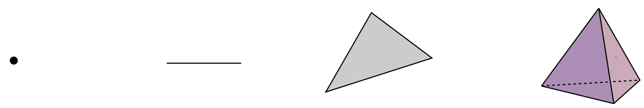

These examples have dimensions 0,1,2, and 3, respectively.

1.3.3 Definition (Simplex). *A simplex σ is the convex hull of a finite affinely independent set A in $\mathbb{R}^d$. The points of A are called the* **vertices** *of σ. The* **dimension** *of σ is $\dim \sigma := |A|-1$. Thus every* **k-simplex** *(k-dimensional simplex) has $k+1$ vertices.*

1.3.4 Definition. *The convex hull of an arbitrary subset of the set of vertices of a simplex σ is a* **face** *of σ. Thus every face is itself a simplex (this is a special case of the definition of a face of a convex polytope).*

The **relative interior** *of a simplex σ arises from σ by removing all faces of dimension smaller than $\dim \sigma$.*

For illustration, we count the faces of a triangle: the whole triangle, three edges, three vertices, and the empty set; altogether we have 8 faces.

Every simplex is a *disjoint* union of the relative interiors of its faces. Thus we get a (closed) triangle as a union of its relative interior (i.e., an open triangle), three open line segments (the edges without their endpoints), and three vertices.

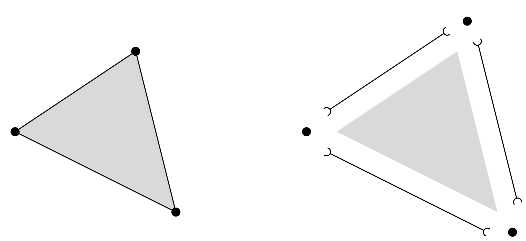

Here are the simple rules of putting simplices together to form a simplicial complex.

1.3.5 Definition. *A nonempty family Δ of simplices is a* **simplicial complex** *if the following two conditions hold:*

(1) *Each face of any simplex $\sigma \in \Delta$ is also a simplex of Δ.*

(2) *The intersection $\sigma_1 \cap \sigma_2$ of any two simplices $\sigma_1, \sigma_2 \in \Delta$ is a face of both σ_1 and σ_2.*

The union of all simplices in a simplicial complex Δ is the **polyhedron** *of Δ and is denoted by $\|\Delta\|$. The* **dimension** *of a simplicial complex is the largest dimension of a face:* $\dim \Delta := \max\{\dim \sigma : \sigma \in \Delta\}$.

In particular, note that every simplicial complex contains the empty set as a face (this is different from some other sources, such as [Mun84] or [Bjö95], where the empty face is excluded!).

The simplicial complex that consists only of the empty simplex is defined to have dimension -1. Zero-dimensional simplicial complexes are just configurations of points, while 1-dimensional simplicial complexes correspond to graphs (represented geometrically with straight edges that do not cross). The following picture shows one 2-dimensional simplicial complex in the plane and two cases of putting simplices together in ways forbidden by the definition of a simplicial complex:

good bad!!!

We are going to restrict ourselves exclusively to *finite* simplicial complexes (with finitely many vertices). From the topological point of view, this is quite a restrictive assumption, since then the polyhedra are only compact spaces and we cannot express, e.g., the space $\mathbb{R}^d$ as the polyhedron of a simplicial complex. But finite simplicial complexes are sufficient for our combinatorial applications and this assumption spares us some trouble (namely, of really discussing too much point set topology).

Support. Just as in the case of a single simplex, the relative interiors of all simplices of a simplicial complex Δ form a partition of $\|\Delta\|$: for each point $\boldsymbol{x} \in \|\Delta\|$ there exists exactly one simplex $\sigma \in \Delta$ containing $\boldsymbol{x}$ in its relative interior. This simplex is denoted by $\operatorname{supp}(\boldsymbol{x})$ and called the *support* of the point $\boldsymbol{x}$.

It may seem obvious at this point that the set of all faces of a simplex forms a simplicial complex—and in fact, this is strongly suggested by our set-up and notation. Still, to be on the safe side, and for further use, we include a proof.

1.3.6 Lemma. *The set of all faces of a simplex is a simplicial complex.*

Proof. Let $V \subset \mathbb{R}^d$ be affinely independent and let $F, G \subseteq V$. We have to show that

$$\operatorname{conv}(F) \cap \operatorname{conv}(G) = \operatorname{conv}(F \cap G),$$

where $\operatorname{conv}(F) \cap \operatorname{conv}(G) \supseteq \operatorname{conv}(F \cap G)$ is trivial. We write $\boldsymbol{x} \in \operatorname{conv}(F) \cap \operatorname{conv}(G)$ as

$$\boldsymbol{x} = \sum_{\boldsymbol{u} \in F} \alpha_{\boldsymbol{u}} \boldsymbol{u} = \sum_{\boldsymbol{v} \in G} \beta_{\boldsymbol{v}} \boldsymbol{v},$$

with $\alpha_{\boldsymbol{u}}, \beta_{\boldsymbol{v}} \geq 0$ and $\sum_{\boldsymbol{u}\in F} \alpha_{\boldsymbol{u}} = 1 = \sum_{\boldsymbol{v}\in G} \beta_{\boldsymbol{v}}$. By subtracting we get

$$\sum_{\boldsymbol{u}\in F\setminus G} \alpha_{\boldsymbol{u}}\boldsymbol{u} - \sum_{\boldsymbol{v}\in G\setminus F} \beta_{\boldsymbol{v}}\boldsymbol{v} + \sum_{\boldsymbol{w}\in F\cap G} (\alpha_{\boldsymbol{w}} - \beta_{\boldsymbol{w}})\boldsymbol{w} = \boldsymbol{0}.$$

The points in $F \cup G$ are affinely independent and thus all the coefficient at the left hand side of this equation must be 0; in particular, $\alpha_{\boldsymbol{w}}, \beta_{\boldsymbol{w}}$ can only be nonzero for $\boldsymbol{w} \in F \cap G$, and thus $\boldsymbol{x} \in \operatorname{conv}(F \cap G)$.

A simplicial complex that is given by an arbitrary n-dimensional simplex and all of its faces will from now on be denoted by σ^n. The n-dimensional simplex itself, as a geometric object, can thus be denoted by $\|\sigma^n\|$.

The notion of subcomplex is defined as everyone would expect:

1.3.7 Definition. *A* **subcomplex** *of a simplicial complex Δ is a subset of Δ that is itself a simplicial complex (that is, it is closed under taking subsets).*

An important example of a subcomplex is the k-*skeleton*of a simplicial complex Δ. It consists of all simplices of Δ of dimension at most k and we denote it by $\Delta^{\leq k}$.

We also use the notation $V(\Delta)$ for the vertex set of Δ.

1.4 Triangulations

Let X be a topological space. A simplicial complex Δ such that $X \cong \|\Delta\|$, if one exists, is called a *triangulation* of X. We give a few examples.

The simplest triangulation of the sphere S^{n-1} is the subcomplex of σ^n obtained by deleting the single n-dimensional simplex (but retaining all of its proper faces). Indeed, the boundary of an n-simplex is homeomorphic to S^{n-1}, as can be seen using the central projection:

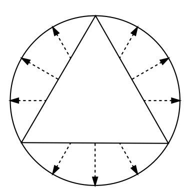

Other triangulations of spheres are obtained from convex polytopes. A convex polytope $P \subset \mathbb{R}^d$ is called *simplicial* if all of its proper faces, (i.e. all faces except possibly for P itself) are simplices. For the familiar 3-dimensional convex polytopes, it means that all the 2-dimensional faces are triangles, as is the case for the regular octahedron or icosahedron. It can be shown without much difficulty that the set of all proper faces of any simplicial polytope P is a simplicial complex. Since $\partial P \cong S^{d-1}$ for every d-dimensional convex polytope P, we obtain various triangulations of the sphere in this way (although, for $d > 3$, by far not all possible triangulations; see Section 5.5!).

Particularly nice and important symmetric triangulations of S^{d-1} are provided by crosspolytopes.

1.4.1 Definition. *The d-dimensional* **crosspolytope** *is the convex hull* $\mathrm{conv}\{\boldsymbol{e}_1, -\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d, -\boldsymbol{e}_d\}$ *of the vectors of the standard orthonormal basis and their negatives:*

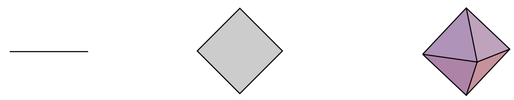

$d = 1$ $\quad$ $d = 2$ $\quad$ $d = 3$

Alternatively, it is the unit ball of the ℓ_1-norm: $\{x \in \mathbb{R}^d : \|\boldsymbol{x}\|_1 \leq 1\}$.

It is not hard to show that a subset $F \subseteq \{\boldsymbol{e}_1, -\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d, -\boldsymbol{e}_d\}$ forms the vertex set of a proper face of the crosspolytope if and only if there is no $i \in [d]$ with both $\boldsymbol{e}_i \in F$ and $-\boldsymbol{e}_i \in F$.

1.4.2 Example (Cube triangulation). The cube $[0, 1]^d$ can be triangulated as follows: Let S_d denote the set of all permutations of $[d]$, and for every $\pi \in S_d$, let $\sigma_\pi = \mathrm{conv}\{\boldsymbol{0}, \boldsymbol{e}_{\pi(1)}, \boldsymbol{e}_{\pi(1)} + \boldsymbol{e}_{\pi(2)}, \ldots, \boldsymbol{e}_{\pi(1)} + \cdots + \boldsymbol{e}_{\pi(d)}\}$. Each σ_π is a d-simplex, and all the σ_π together plus all of their faces form a triangulation of $[0, 1]^d$ (we leave the verification as Exercise 3).

Notes. To construct "suitable" triangulations of given geometric shapes is a major topic in many fields of Applied Mathematics, such as Numerical Analysis and Computer Aided Design (CAD).

In contemporary algebraic topology, simplicial complexes are often considered old-fashioned. Spaces can be usually described much more economically if we allow for more general ways of gluing the basic building blocks together than is permitted in simplicial complexes. For example, the torus (the surface of a tire-tube) can be produced by a suitable gluing of the edges of a single square in $\mathbb{R}^3$,

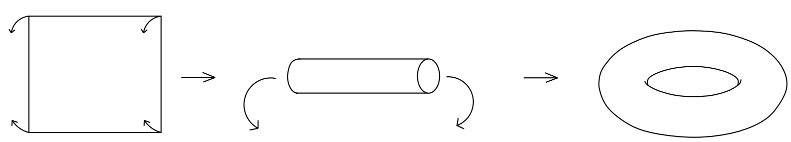

while a triangulation of the torus requires quite a number of simplices (at least 14 triangles, in fact). Moreover, there are quite "reasonable" spaces (4-dimensional manifolds) which cannot be triangulated at all, while they can be obtained using more general ways of gluing.

However, these more general ways of building spaces, most notably the CW-complexes (briefly discussed in Section 4.4), do not admit as direct combinatorial interpretation as simplicial complexes do.

Exercises

1. Draw a triangulation of a torus. Use as few simplices as you can.

2. (a) Prove the claim about the faces of the crosspolytope below Definition 1.4.1 (use the definition of a polytope face mentioned in the Preliminaries).

   (b) Count the number of faces of each dimension.

3. This refers to the cube triangulation in Example 1.4.2.

   (a) Check that each simplex σ_π is d-dimensional and can be written as $\sigma_\pi = \{\boldsymbol{x} \in [0,1]^d : x_{\pi(d)} \leq x_{\pi(d-1)} \leq \cdots \leq x_{\pi(1)}\}$. Conclude that $\bigcup_{\pi \in S_n} \sigma_\pi = [0,1]^d$.

   (b) Let $\preceq$ be a *linear quasiordering* of $[d]$, i.e. a transitive relation in which every two numbers are comparable, $i \preceq j$ or $j \preceq i$ (but it may happen that both $i \preceq j$ and $j \preceq i$ even if $i \neq j$). Define $\sigma_\preceq := \{\boldsymbol{x} \in [0,1]^d : x_i \leq x_j \text{ whenever } i \preceq j\}$. Check that $\sigma_\preceq$ is a simplex, determine its dimension (in terms of $\preceq$), and describe its vertices.

   (c) Show that the intersection $\sigma_{\preceq_1} \cap \sigma_{\preceq_2}$ again of the form $\sigma_\preceq$ for a suitable linear quasiordering $\preceq$. How do we obtain $\preceq$ from $\preceq_1$ and $\preceq_2$?

   (d) What are the faces of σ_π? Verify that the σ_π and their faces form a simplicial complex.

   (e) Show that the copies of the triangulation in Example 1.4.2 translated by each integer vector in $\{0,1,\ldots,n-1\}$ form a triangulation of $[0,n]^d$.

1.5 Abstract simplicial complexes

We introduce a combinatorial notion which later on turns to be equivalent to a geometric simplicial complex.

1.5.1 Definition. *An* **abstract simplicial complex** *is a pair* (V, K), *where* V *is a set and* $\mathsf{K} \subseteq 2^V$ *is a hereditary system of subsets of* V*; that is, we require that* $F \in \mathsf{K}$ *and* $G \subseteq F$ *imply* $G \in \mathsf{K}$. *The sets in* K *are called (abstract) simplices. Further we define the* **dimension** $\dim(\mathsf{K}) := \max\{|F|-1 : F \in \mathsf{K}\}$.

Usually we may assume that $V = \bigcup \mathsf{K}$; thus it suffices to write K instead of (V, K), where V is understood to equal $\bigcup \mathsf{K}$.

Each geometric simplicial complex Δ determines an abstract simplicial complex. The points of the abstract simplicial complex are all vertices of the simplices of Δ, so we set $V := V(\Delta)$, and the sets in the abstract simplicial complex are just the vertex sets of the simplices of Δ. The set system (V, K) obtained in this way is clearly an abstract simplicial complex.

In this situation, we call Δ a *geometric realization* of K, and the polyhedron of Δ is also referred to as a *polyhedron of* K (soon we will see that a polyhedron of K is unique up to homeomorphism).

It is easy to see that any abstract simplicial complex (V, K) with V finite (which we always assume) has a geometric realization. Let $n := |V|-1$ and let us identify V with the vertex set of an n-dimensional simplex $\sigma^n \subset \mathbb{R}^n$. We define a subcomplex Δ of σ^n: $\Delta = \{\mathrm{conv}(F) : F \in \mathsf{K}\}$. Quite obviously, this is a geometric simplicial complex and its associated abstract simplicial complex is just K. So every simplicial complex on $n+1$ vertices can be realized in $\mathbb{R}^n$ (later on, we will prove a much sharper result).

Now we show that the geometric realization is unique up to homeomorphism. At this occasion, we also introduce the important notion of a simplicial mapping.

1.5.2 Definition. *Let K and L be two abstract simplicial complexes. A* **simplicial mapping** *of K into L is a mapping $f: V(\mathsf{K}) \to V(\mathsf{L})$ that maps simplices to simplices, i.e. such that $f(F) \in \mathsf{L}$ whenever $F \in \mathsf{K}$.*

A bijective simplicial mapping whose inverse mapping is also simplicial is called an **isomorphism** *of abstract simplicial complexes.*

Isomorphic abstract simplicial complexes are thus "the same" set systems, they only differ in the names of the vertices. In the sequel, we won't usually distinguish among isomorphic simplicial complexes.

We also note that for an arbitrary simplicial mapping, a k-simplex in K can be mapped to a simplex of L of any dimension $\ell \leq k$.

To each simplicial mapping f of simplicial complexes, we are going to associate a continuous mapping $\|f\|$ of their polyhedra. Namely, we extend f affinely on each simplex. To state this precisely, we first note that if $\sigma \subset \mathbb{R}^d$ is a k-simplex with vertices $\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$, then each point $\boldsymbol{x} \in \sigma$ can be uniquely written as a convex combination $\boldsymbol{x} = \sum_{i=0}^k \alpha_i \boldsymbol{v}_i$, where $\alpha_0, \ldots, \alpha_k \geq 0$ and $\sum_{i=0}^k \alpha_i = 1$. Indeed, at least one such convex combination exists because $\boldsymbol{x} \in \operatorname{conv}\{\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k\}$, and if there were two distinct convex combinations equal to $\boldsymbol{x}$, we would get a contradiction to the affine independence of $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k$ by subtracting them.

1.5.3 Definition. *Let Δ_1 and Δ_2 be geometric simplicial complexes, let K_1 and K_2 be their associated abstract simplicial complexes, and let $f: V(\mathsf{K}_1) \to V(\mathsf{K}_2)$ be a simplicial mapping of K_1 into K_2. We define the mapping $\|f\|: \|\Delta_1\| \to \|\Delta_2\|$ by extending f affinely to the relative interiors of the simplices of Δ_1, as follows: if $\sigma = \operatorname{supp}(\boldsymbol{x}) \in \Delta_1$ is the support of $\boldsymbol{x}$, the vertices of σ are $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k$, and $\boldsymbol{x} = \sum_{i=0}^k \alpha_i \boldsymbol{v}_i$, we put $\|f\|(x) = \sum_{i=0}^k \alpha_i f(\boldsymbol{v}_i)$.*

First we note that the mapping $\|f\|$ is well-defined, because $\{f(\boldsymbol{v}_0), \ldots, f(\boldsymbol{v}_k)\}$ is always the vertex set of a simplex in Δ_2 as f is simplicial. With some more effort, one can check the following proposition, whose proof we omit.

1.5.4 Proposition. *For every simplicial mapping f as in Definition 1.5.3, $\|f\|$ is a continuous map $\|\Delta_1\| \to \|\Delta_2\|$. If f is injective or surjective, then $\|f\|$ has the same property and if f is an isomorphism, then $\|f\|$ is a homeomorphism.*

In particular, this proposition shows that each (finite) abstract simplicial complex (V, K) defines a topological space uniquely up to homeomorphism.

Convention. In the sequel, a simplicial complex will formally be understood as an abstract simplicial complex (i.e. it will be a set system as a mathematical object). But we will speak of a polyhedron $\|\mathsf{K}\|$ for an abstract simplicial complex K (which is well-defined up to homeomorphism in view of Proposition 1.5.4) and even use topological notions such as "K is contractible," instead of "$\|K\|$ is contractible."

1.6 Dimension of geometric realizations

Here is the promised sharper result about realizability of d-dimensional simplicial complexes.

1.6.1 Theorem (Geometric realization theorem). *Every finite d-dimensional simplicial complex K has a geometric realization in $\mathbb{R}^{2d+1}$.*

For $d = 1$, the theorem says that every graph can be represented in $\mathbb{R}^3$, with edges being straight segments. The dimension 3 is the smallest possible in general since there are non-planar graphs. We will later show that $2d+1$ is the smallest possible dimension for all d; see the Van Kampen–Flores theorem 5.5.2. Of course, this applies only in the worst case, since there are many d-dimensional simplicial complexes which can be realized in dimensions lower than $2d+1$ (say the d-simplex).

In the proof of Theorem 1.6.1, we use the following sufficient condition for a geometric realization.

1.6.2 Lemma. *If K is a simplicial complex and $f\colon V(\mathsf{K}) \to \mathbb{R}^d$ is an injective map such that $f(F \cup G)$ is affinely independent for all $F, G \in \mathsf{K}$, then the assignment*

$$F \longmapsto \sigma_F := \mathrm{conv}(f(F))$$

provides a geometric realization of K in $\mathbb{R}^d$.

Proof. If $f(F \cup G)$ is affinely independent, then σ_F and σ_G are two faces of the simplex with the vertex set $f(F \cup G)$, and we are done by Lemma 1.3.6.

A suitable placement of vertices can be defined using the moment curve. Later on, we will meet this useful curve several more times.

1.6.3 Definition. *The curve $\{\gamma(t) : t \in \mathbb{R}\}$ given by $\gamma(t) := (t, t^2, \ldots, t^d)$ is the* **moment curve** *in $\mathbb{R}^d$.*

The following lemma expresses a key property of the moment curve (any curve with this property would do in the sequel). It is a little stronger than needed here.

1.6.4 Lemma. *No hyperplane intersects the moment curve γ in $\mathbb{R}^d$ in more than d points. Consequently, every set of $d+1$ distinct points on γ is affinely independent. Moreover, if γ intersects a hyperplane h at d distinct points, then it crosses h from one side to the other at each intersection.*

Proof. A hyperplane h has an equation $\langle \boldsymbol{a}, \boldsymbol{x}\rangle = b$ with $\boldsymbol{a} \neq \boldsymbol{0}$. If a point $\gamma(t)$ lies in h, then we have $a_1 t + a_2 t^2 + \cdots + a_d t^d = b$. This means that the values of t corresponding to intersections with h are the roots of the nonzero polynomial $p(t) = (\sum_{i=1}^{d} a_i t^i) - b$ of degree at most d. Such a $p(t)$ has at most d roots, and so there are no more than d intersections.

If there are d distinct intersections, then $p(t)$ has d distinct roots, which must be all simple. Therefore, $p(t)$ changes sign at each root, and this means that γ passes from one open halfspace defined by h to the other at each intersection.

Proof of Theorem 1.6.1. We choose a map $f\colon V(\mathsf{K}) \to \mathbb{R}^{2d+1}$ such that the vertices of K are assigned distinct points on the moment curve in $\mathbb{R}^{2d+1}$. Then for $F, G \in \mathsf{K}$ we have $|F \cup G| \leq (d+1) + (d+1) = 2d+2$, and thus by Lemma 1.6.4 the corresponding points in $f(F \cup G)$ are affinely independent. Hence we are done by Lemma 1.6.2.

Exercises

1. The *chessboard complex* $\text{♖}_{m,n}$ has the squares of the $m \times n$ chessboard as vertices, and simplices are all subsets of squares such that no two squares lie in the same row or column (so if we place rooks to these squares they do not threaten one another). Describe the "geometric shape" of $\|\text{♖}_{3,4}\|$.

1.7 Simplicial complexes and posets

We recall that a *partially ordered set*, or *poset* for short, is a pair $(P, \preceq)$, where P is a set and $\preceq$ is a binary relation on P that is reflexive ($x \preceq x$), transitive ($x \preceq y$ and $y \preceq z$ implies $x \preceq z$), and weakly antisymmetric ($x \preceq y$ and $y \preceq x$ implies $x = y$). Similar to topological spaces, $\preceq$ is sometimes omitted from the notation.

As we will see, there is a correspondence between (finite) simplicial complexes and (finite) posets. It is not quite one-to-one but each poset is assigned a unique topological space, up to homeomorphism.

1.7.1 Definition. *The* **order complex** *of a poset P is the simplicial complex $\Delta(P)$, whose vertices are the elements of P and whose simplices are all chains (i.e. linearly ordered subsets, of the form $\{x_1, x_2, \ldots, x_k\}$, $x_1 \prec x_2 \prec \cdots \prec x_k$) in P.*

The **face poset** *of a simplicial complex K is the poset $P(\mathsf{K})$, which is the set of all* **nonempty** *simplices of K, ordered by inclusion.*

For example, the simplicial complex

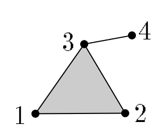

has the face poset

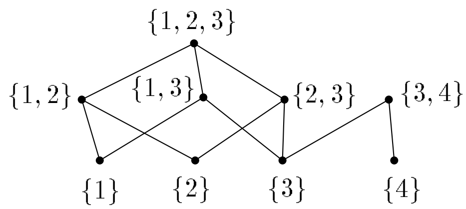

(this is the *Hasse diagram* of the poset, where each element is connected to its immediate predecessors and immediate successors, with the predecessors lying below it and the successors above it). Here is the order complex of this poset, together with a Meadow Saffron (or also Autumn Crocus, *Colchicum autumnale L.*) as an extra bonus:

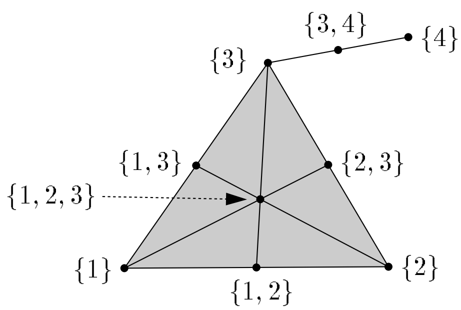

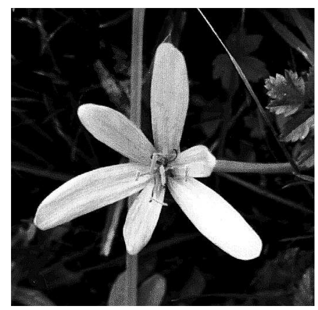

The operation we just did on the original simplicial complex, namely passing to the face poset and then to its order complex, is very important and has a name:

1.7.2 Definition. *For a simplicial complex* K*, the simplicial complex* $\mathrm{sd}(\mathsf{K}) := \Delta(P(\mathsf{K}))$ *is called the* **(first) barycentric subdivision** *of* K.

More explicitly, the vertices of $\mathrm{sd}(\mathsf{K})$ are the nonempty simplices of K and the simplices of $\mathrm{sd}(\mathsf{K})$ are chains of simplices of K ordered by inclusion.

Given a geometric realization of K, we can place the vertex of $\mathrm{sd}(\mathsf{K})$ corresponding to a simplex σ to the center of gravity (barycenter) of σ, as we did in the above picture. It turns out that, as the picture suggests, $\|\mathrm{sd}(\mathsf{K})\|$ is always (canonically) homeomorphic to $\|\mathsf{K}\|$. It suffices to prove this for K being (the simplicial complex of) a simplex; this is not very difficult and we leave it to reader's diligence.

In algebraic topology, mainly in the earlier days, iterated barycentric subdivision was used for constructing arbitrarily fine triangulations of a given polyhedron. In the applications in this text, we will mainly encounter barycentric subdivision in its combinatorial meaning, in connection with posets.

Monotone maps and simplicial maps. Let $(P_1, \preceq_1)$ and $(P_2, \preceq_2)$ be posets. A mapping $f\colon P_1 \to P_2$ is called *monotone* if $x \preceq_1 y$ implies $f(x) \preceq_2 f(y)$. We have the following simple but useful

1.7.3 Proposition. *Every monotone mapping* $f\colon P_1 \to P_2$ *between posets is also a simplicial mapping* $\Delta(P_1) \to \Delta(P_2)$ *between their order complexes.*

We again leave the very easy verification to the reader.

1.7.4 Corollary. *Let K_1 and K_2 be simplicial complexes. Consider an arbitrary mapping f which assigns to each simplex $F \in \mathsf{K}_1$ a simplex $f(F) \in \mathsf{K}_2$ (f is not necessarily induced by a mapping of vertices!), and suppose that if $F' \subseteq F$, then also $f(F') \subseteq f(F)$. Then f can be regarded as a simplicial mapping $f: \mathrm{sd}(\mathsf{K}_1) \to \mathrm{sd}(\mathsf{K}_2)$, and so it induces a continuous map $\|f\|: \|\mathsf{K}_1\| \to \|\mathsf{K}_2\|$.*

Notes. The order complex $\Delta(P)$ is an instance of a more general construction of a *classifying space*; see e.g. [Hat01, Chapter 2].

Let us mention a result somewhat similar to the geometric realization theorem (Theorem 1.6.1), which provides an upper bound on the dimension necessary for embedding a given simplicial complex. First we recall the notion of *Dushnik–Miller dimension* (or *order dimension*) of a poset. As is easy to check, if $(P, \preceq)$ is a finite poset, there exist linear orderings $\leq_1, \leq_2, \ldots, \leq_k$ such that $x \preceq y$ iff $x \leq_i y$ for all $i \in [k]$ (in other words, $\preceq = \bigcap_{i=1}^k \leq_i$). The smallest possible k for such a representation of $\preceq$ by linear orderings is the Dushnik–Miller dimension $\dim(P, \preceq)$. Ossona de Mendez [Oss99] proved, using so-called Scarf's construction, that every finite simplicial complex K can be geometrically realized in $\mathbb{R}^{d-1}$ with $d = \dim(P(\Delta))$. For a proof, let $\leq_1, \ldots, \leq_d$ be linear orderings of K witnessing $\dim(P(\mathsf{K})) = d$. We restrict the orderings $\leq_i$ to the set $V := V(\mathsf{K})$ (the vertices are also simplices of K) and let φ_i be the injective map $V \to [n]$, $n = |V|$, that is monotone with respect to $\leq_i$ (that is, $u <_i v$ iff $\varphi_i(u) < \varphi_i(v)$ for every $u, v \in V$). Define $f_0: V \to \mathbb{R}^d$ by $f_0(v) = ((d{+}1)^{\varphi_1(v)}, (d{+}1)^{\varphi_2(v)}, \ldots, (d{+}1)^{\varphi_d(v)})$ and finally let $f(v)$ be the projection of $f_0(v)$ from $\mathbf{0}$ on the hyperplane $\sum_{i=1}^d x_i = 1$. Then it can be shown that f satisfies the condition of Lemma 1.6.2 and thus provides a realization of K in $\mathbb{R}^{d-1}$.

A converse of this theorem is known for $d = 3$: if we regard a graph G as a 1-dimensional simplicial complex, then the dimension of the face poset is at most 3 if and only if G is planar [Sch89]; also see [BT93], [BT97] for related results.

Exercises

1. Prove that a simplex is homeomorphic to its barycentric subdivision (a rigorous proof takes some work!).

2. Prove Proposition 1.7.3 and Corollary 1.7.4.

2

The Borsuk–Ulam Theorem

The Borsuk–Ulam theorem is one of the most useful tools offered by elementary algebraic topology to the outside world. Here are four reasons why this is such a great theorem: there are

(1) several different equivalent versions,

(2) many different proofs,

(3) a host of extensions and generalizations, and

(4) numerous interesting applications.

As for (1), below we give eight different but equivalent versions, all of them very useful. They include all three versions from Borsuk's original 1933 paper [Bor33].

As for (2), there are several proofs of the Borsuk–Ulam theorem that can be labeled as completely elementary, requiring just undergraduate mathematics and no algebraic topology. On the other hand, most of the textbooks on algebraic topology, even the friendliest ones, usually place a proof of the Borsuk–Ulam theorem well beyond page 100. Some of them use just basic homology theory, others rely on properties of the cohomology ring, but in any case, significant apparatus has to be mastered for really understanding such proofs. From a "higher" point of view, it can be argued that these proofs are more conceptual and go to the heart of the matter, and thus they are preferable to the "ad hoc" elementary proofs. But this point of view can only be appreciated by someone for whom the necessary machinery is as natural as breathing. Since not everyone, especially in combinatorics and computer science, belongs to this lucky group, we present two "old-fashioned" elementary proofs. The one in Section 2.2, a so-called homotopy extension argument, is geometric and very intuitive. The other, in Section 2.3, resembling the proof of Brouwer's theorem via the Sperner lemma, derives the Borsuk–Ulam theorem from a purely combinatorial statement called Tucker's lemma.

As for (3), we will examine various generalizations and strengthenings later; much more can be found in Steinlein's surveys [Ste85], [Ste93] and in the sources he quotes.

Finally, as for applications (4), just wait and see.

2.1 The Borsuk–Ulam theorem in various guises

One of the versions of the Borsuk–Ulam theorem, the one that is perhaps the easiest to remember, states that *for any continuous mapping* $f\colon S^n \to \mathbb{R}^n$, *there exists a point* $\boldsymbol{x} \in S^n$ *such that* $f(\boldsymbol{x}) = f(-\boldsymbol{x})$. This is schematically indicated below:

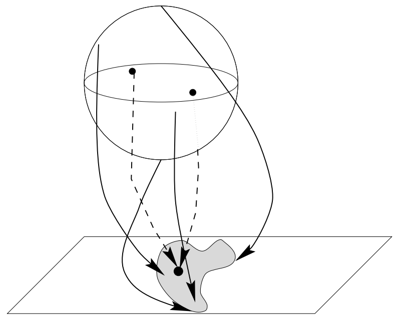

A popular interpretation found in almost every textbook says that at any given time, there are two antipodal places on the Earth that have the same temperature and, at the same time, identical air pressure (here $n = 2$).[1]

It is instructive to compare this with the Brouwer fixed point theorem, which says that every continuous mapping $f\colon B^n \to B^n$ has a fixed point: $f(\boldsymbol{x}) = \boldsymbol{x}$ for some $\boldsymbol{x} \in B^n$. The statement of the Borsuk–Ulam theorem sounds similar (and, actually, it easily implies the Brouwer theorem), but it involves an extra ingredient besides the topology of the considered spaces: certain *symmetry* of these spaces, namely the symmetry given by the mapping $\boldsymbol{x} \mapsto -\boldsymbol{x}$ (which is often called the *antipodality* on S^n and on $\mathbb{R}^n$, respectively).

Here is Borsuk's original version of the Borsuk–Ulam theorem:

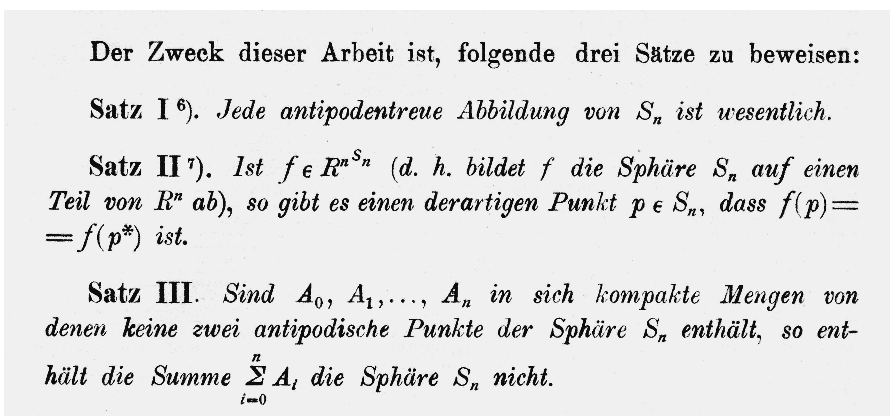

Der Zweck dieser Arbeit ist, folgende drei Sätze zu beweisen:

Satz I [6]). *Jede antipodentreue Abbildung von* S_n *ist wesentlich.*

Satz II [7]). *Ist* $f \in {R^n}^{S_n}$ *(d. h. bildet* f *die Sphäre* S_n *auf einen Teil von* R^n *ab), so gibt es einen derartigen Punkt* $p \in S_n$, *dass* $f(p) = f(p^*)$ *ist.*

Satz III. *Sind* $A_0, A_1, \ldots, A_n$ *in sich kompakte Mengen von denen keine zwei antipodische Punkte der Sphäre* S_n *enthält, so enthält die Summe* $\sum_{i=0}^{n} A_i$ *die Sphäre* S_n *nicht.*

Here are the promised many equivalent versions, in English; the statements most significant for us are those with boldface numbers.

2.1.1 Theorem (Borsuk–Ulam theorem). *For all* $n \geq 0$, *the following statements are equivalent, and true:*

[1] Although anyone who has ever touched a griddle-hot stove knows that the temperature need not be continuous.

(1.1) (Borsuk [Bor33, Satz II][2]) *For every continuous mapping* $f\colon S^n \to \mathbb{R}^n$ *there exists a point* $\boldsymbol{x} \in S^n$ *with* $f(\boldsymbol{x}) = f(-\boldsymbol{x})$.

(1.2) *For every* **antipodal mapping** $f\colon S^n \to \mathbb{R}^n$ *(that is, f is continuous and $f(\boldsymbol{x}) = -f(-\boldsymbol{x})$ for all $\boldsymbol{x} \in S^n$) there exists a point $\boldsymbol{x} \in S^n$ satisfying $f(\boldsymbol{x}) = \mathbf{0}$.*

(1.3) *Let $g\colon B^n \to \mathbb{R}^n$ be a continuous map that satisfies $g(-\boldsymbol{x}) = -g(\boldsymbol{x})$ for all $\boldsymbol{x} \in S^{n-1}$; that is, it is antipodal on the boundary. Then there is a point $\boldsymbol{x}^* \in B^n$ with $g(\boldsymbol{x}^*) = \mathbf{0}$.*

(2.1) *There is no antipodal mapping $f\colon S^n \to S^{n-1}$.*

(2.2) *An antipodal map $f\colon S^{n-1} \to S^{n-1}$ cannot be nullhomotopic.*

(2.3) (Borsuk [Bor33, Satz I][3]) *If $f\colon S^{n-1} \to S^{n-1}$ is antipodal, then every map $g\colon S^{n-1} \to S^{n-1}$ that is homotopic to f is surjective.*

(3.1) (Lusternik & Schnirelmann [LS30], Borsuk [Bor33, Satz III]) *For any cover $B_1, \ldots, B_{n+1}$ of the sphere S^n by $n+1$ closed sets, there is at least one set containing a pair of antipodal points (that is, $B_i \cap (-B_i) \neq \emptyset$).*

(3.2) *For any cover $A_1, \ldots, A_{n+1}$ of the sphere S^n by $n+1$ open sets, there is at least one set containing a pair of antipodal points.*

While proving any of the versions of the Borsuk–Ulam theorem is not easy, at least without some technical apparatus, checking the equivalence of all the statements is not so hard. Deriving at least some of the equivalences before reading further is a very good way of getting a feeling for the theorem. Here we begin with the boldface statements.

Equivalence of (1.1), (1.2), and (2.1).
(1.1) $\Longrightarrow$ (1.2) is clear.
(1.2) $\Longrightarrow$ (1.1) We convert f into an antipodal mapping by setting $g(\boldsymbol{x}) := f(\boldsymbol{x}) - f(-\boldsymbol{x})$.
(1.2) $\Longrightarrow$ (2.1) An antipodal map $S^n \to S^{n-1}$ is also a nowhere zero antipodal mapping $S^n \to \mathbb{R}^n$.
(2.1) $\Longrightarrow$ (1.2) Assume that $f\colon S^n \to \mathbb{R}^n$ is a continuous nowhere zero antipodal mapping. Then the antipodal mapping $g\colon S^n \to S^{n-1}$ given by $g(\boldsymbol{x}) := f(\boldsymbol{x})/\|f(\boldsymbol{x})\|$ contradicts (2.1).

Equivalence of (1.3) with (1.2) is easy once we observe that the projection $\pi\colon (x_1, \ldots, x_{n+1}) \mapsto (x_1, \ldots, x_n)$ is a homeomorphism of the upper hemisphere U of S^n with B^n:

[2] Borsuk's footnote at this theorem reads: "This theorem was posed as a conjecture by St. Ulam."

[3] Borsuk's footnote at this point shows his suffering from the fact that there are many proofs: "Mr. H. Hopf, whom I informed about Theorem I, noted for me in a letter three other shorter proofs of this theorem. But since these proofs are founded on deep results in the theory of the mapping degree and my proof is in essence completely elementary, I think that its publication is not superfluous. [...]"

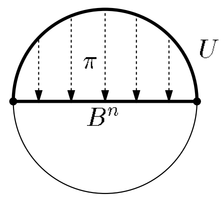

An antipodal mapping $f: S^n \to \mathbb{R}^n$ as in (1.2) thus yields a mapping $g: B^n \to \mathbb{R}^n$ antipodal on ∂B^n by $g(\boldsymbol{x}) = f(\pi^{-1}(\boldsymbol{x}))$. Conversely, for a $g: B^n \to \mathbb{R}^n$ as in (1.3) we can define $f(\boldsymbol{x}) = g(\pi(\boldsymbol{x}))$ and $f(-\boldsymbol{x}) = -g(\pi(\boldsymbol{x}))$ for $\boldsymbol{x} \in U$. This specifies f on the whole S^n, it is consistent because g is antipodal on the equator of S^n, and the resulting f is continuous since it is continuous on both the closed hemispheres (see Exercise 1.1.2).

Equivalence with the Lusternik–Schnirelmann theorem (3.1), (3.2). (1.1) $\Longrightarrow$ (3.1) For a closed cover $B_1, \ldots, B_{n+1}$ we define a continuous mapping $f: S^n \to \mathbb{R}^n$ by $f(\boldsymbol{x}) := (\operatorname{dist}(\boldsymbol{x}, B_1), \ldots, \operatorname{dist}(\boldsymbol{x}, B_n))$ and we consider a point $\boldsymbol{x} \in S^n$ with $f(\boldsymbol{x}) = f(-\boldsymbol{x}) = \boldsymbol{y}$, which exists by (1.1). If the ith coordinate of the point $\boldsymbol{y}$ is 0, then both $\boldsymbol{x}$ and $-\boldsymbol{x}$ are in B_i. If all coordinates of $\boldsymbol{y}$ are nonzero, then both $\boldsymbol{x}$ and $-\boldsymbol{x}$ lie in B_{n+1}.
(3.1) $\Longrightarrow$ (2.1) We need an *auxiliary result*: There exists a covering of S^{n-1} by closed sets $B_1, \ldots, B_{n+1}$ such that no B_i contains a pair of antipodal points (to see this, we can use the projection of the faces of a simplex that has the origin in its interior). Then if a continuous antipodal mapping $f: S^d \to S^{n-1}$ with $d \geq n$ existed, then the sets $f^{-1}(B_1), \ldots, f^{-1}(B_{n+1})$ would contradict (3.1).
(3.1) $\Longrightarrow$ (3.2) follows from the fact that for every open cover $A_1, \ldots, A_{n+1}$ there exists a closed cover $B_1, \ldots, B_{n+1}$ satisfying $B_i \subset A_i$ for $i = 1, \ldots, n+1$: for each point $\boldsymbol{x}$ of the sphere choose its open neighborhood $U_{\boldsymbol{x}}$ whose closure is contained in some A_i, and apply the compactness of the sphere.
(3.2) $\Longrightarrow$ (3.1) follows from the fact that each set of a closed cover $B_1, \ldots, B_{n+1}$ can be wrapped in an open set $A_i^{\varepsilon} = \{\boldsymbol{x} \in S^n : \operatorname{dist}(\boldsymbol{x}, B_i) < \varepsilon\}$. We let $\varepsilon \to 0$ and we use the compactness of the sphere. Taking twice a suitable infinite subsequence, we first obtain an infinite sequence of points $\boldsymbol{x}^0, \boldsymbol{x}^1, \boldsymbol{x}^2, \ldots$ in S^n such that $\operatorname{dist}(\boldsymbol{x}^i, B_j), \operatorname{dist}(-\boldsymbol{x}^i, B_j) \to 0$ for $i \to \infty$ some *fixed* j, and then a *convergent* subsequence. The limit point of this sequence is in B_j since B_j is closed, and this provides the required antipodal pair in B_j.

Finally, we leave the equivalence with the "homotopic" statements (2.2) and (2.3) to the exercises.

Proof of the Brouwer fixed point theorem from the Borsuk–Ulam theorem (2.2). Suppose that $f: B^n \to B^n$ is continuous and has no fixed point. By a well-known construction, we show the existence of a continuous map $g: B^n \to S^{n-1}$ whose restriction to S^{n-1} is the identity map (such a g is called a *retraction* of B^n to S^{n-1}). We define $g(\boldsymbol{x})$ as the point in which the ray originating in $f(\boldsymbol{x})$ and going through $\boldsymbol{x}$ intersects S^{n-1}. This g, considered as a mapping $B^n \to \mathbb{R}^n$, contradicts version (1.3) of the Borsuk–Ulam theorem.

Notes. The earliest reference for what is now commonly called the Borsuk–Ulam theorem is probably Lusternik & Schnirelmann [LS30] from 1930 (the covering version (3.1)). Borsuk's paper [Bor33] is from 1933. Since then, hundreds of papers with various new proofs, variations of old proofs, generalizations, and applications, have appeared; the most comprehensive survey known to us, Steinlein [Ste85] from 1985, lists nearly 500 items in the bibliography.

Types of proofs. In the numerous published proofs of the Borsuk–Ulam theorem, one can distinguish several basic approaches (as is done in [Ste85]). Some of these types will be treated in this text; for the others, we outline the main ideas here and give references, mostly to recent textbooks.

In *degree-theoretical proofs*, one shows that a continuous antipodal mapping $f\colon S^n \to S^n$ has odd degree; this implies that it cannot be nullhomotopic (version (2.1)) since a nullhomotopic map has degree 0. Here the degree can be defined homologically, as the number d such that the homomorphism $f_*\colon H_n(S^n, \mathbb{Z}) \cong \mathbb{Z} \to H_n(S^n, \mathbb{Z}) \cong \mathbb{Z}$ induced by f in the nth homology acts as the multiplication by d (see Dodson and Parker [DP97, Sec. 4.3.2] for such a proof). Another, more universal, definition of degree uses algebraic counting of the roots x of $f(x) = y$ at a "generic" image point y. In particular, for the purposes of the Borsuk–Ulam theorem, it suffices to define the degree modulo 2, and then it is congruent mod 2 to the number of preimages of a generic point y. A proof using the degree of a smooth map is sketched in [Bre93, p. 253]. In the degree-theoretical approach, one has to approximate the arbitrary antipodal map by a suitable nice (simplicial, or smooth) map so that the degree is well-defined. A related method uses the *Lefschetz number*; see Section 6.2. A proof using rudimentary *Smith theory* can be found in [Bre93, Sec. 20].

A *proof using the cohomology ring* considers the map $g\colon \mathbb{R}\mathrm{P}^n \to \mathbb{R}\mathrm{P}^m$ induced by an antipodal $f\colon S^n \to S^m$, and shows that the corresponding homomorphism $g^*\colon H^*(\mathbb{R}\mathrm{P}^m, \mathbb{Z}_2) \to H^*(\mathbb{R}\mathrm{P}^n, \mathbb{Z}_2)$ of the cohomology rings carries a generator α of $H^1(\mathbb{R}\mathrm{P}^m, \mathbb{Z}_2)$ to a generator β of $H^1(\mathbb{R}\mathrm{P}^n, \mathbb{Z}_2)$. This is impossible if $m + 1 \leq n$, since then α^{m+1} is trivial while β^n is nontrivial. See, for example, [Mun84, p. 403] or [Bre93, p. 362].

A proof by a *homotopy extension argument* will be discussed in Section 2.2, and a representative of the family of *combinatorial proofs* in Section 2.3. *Algebraic proofs* were given in [Kne82] and [AP83].

As for *applications* of the Borsuk–Ulam theorem, we will cover some number in the subsequent sections. For a multitude of others, we refer to the surveys [Ste85], [Ste93]. The papers [Bár93] and [Alo88] give nice overviews of combinatorial applications; most of these are included in this text.

A very broad field of applications, which we will neglect entirely, are existence results for solutions of nonlinear *partial differential equations* and *integral equations*. Also in *functional analysis* and *geometry of Banach spaces*, Borsuk–Ulam type results play an important role.

Bourgin–Yang type theorems are generalizations of the Borsuk–Ulam theorem of the following sort. For any continuous map $f\colon S^n \to \mathbb{R}^m$, the *coincidence set* $\{\boldsymbol{x} \in S^n : f(\boldsymbol{x}) = f(-\boldsymbol{x})\}$ has to be not only nonempty (as Borsuk–Ulam asserts), but even "large" if $m < n$; for example, it has dimension at least $n - m$. Such results have been used in proving various geometric statements. We will mention a little more about this in the notes to Section 5.2.

A beautiful combinatorial application that we will not discuss in detail (for space reasons, and also because the original account is nicely readable) concern *linkless embeddings* of graphs in $\mathbb{R}^3$. Any finite graph G, regarded as a 1-dimensional finite simplicial complex, can be realized in $\mathbb{R}^3$. Such a realization is called *linkless* if any two vertex-disjoint circuits in G form two unlinked closed curves in the realization. Here two curves $\alpha, \beta \subset \mathbb{R}^3$ (each homeomorphic to S^1) are unlinked if they are equivalent to two isometric copies α', β' of S^1 in $\mathbb{R}^3$ lying far from one another, and the equivalence means that there is a homeomorphism $\varphi\colon \mathbb{R}^3 \to \mathbb{R}^3$ such that $\varphi(\alpha \cup \beta) = \alpha' \cup \beta'$ (these are notions from knot theory; see e.g. Rolfsen [Rol90] for more information).

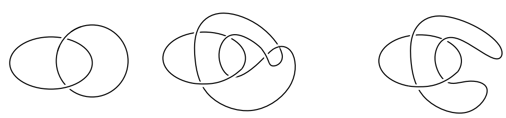

linked linked unlinked

Lovász and Schrijver [LS98], building on previous work by Robertson, Seymour, and Thomas, proved that graphs possessing a linkless embedding into $\mathbb{R}^3$ are exactly those for which a numerical parameter, called the *Colin de Verdière number*, is at most 4. The definition of this parameter, using spectra of certain matrices, is not very intuitive at first sight (and we do not reproduce it; see, for instance, the book [Col99]). The graph-theoretical significance of the Colin de Verdière number looks almost miraculous: besides the incredible result about linkless embeddings, it is known, for instance, that the class of graphs having this parameter at most 3 are exactly all planar graphs! In the Lovász–Schrijver proof, the Borsuk–Ulam theorem is used for establishing the following: Given any "generic" embedding of the 1-skeleton of a 5-dimensional convex polytope P into $\mathbb{R}^3$, there are two antipodal 2-dimensional faces of P such that the images of the boundaries of these two faces are linked (in fact, they have a nonzero linking number, which is stronger than being linked—the curves in the left picture above satisfy this while those in the middle picture do not). Thus, for example, the complete graph K_6 is not linklessly embeddable. (More generally, a generic embedding of the $(d-1)$-skeleton of a $(2d+1)$-polytope into $\mathbb{R}^d$ links the boundaries of two antipodal d-faces.)

The paper [Bor33] containing the Borsuk–Ulam theorem also states the so-called *Borsuk's conjecture* [Bor33]. The Lusternik–Schnirelmann theorem (about covering S^n by $n+1$ closed sets) can be restated as follows: For every closed cover of S^{n-1} by at most n sets, one of the sets has diameter 2, i.e. the same as the diameter of S^n itself. On the other hand, there are $n+1$ sets of diameter < 2 covering S^n. Borsuk asked if any bounded set $X \subset \mathbb{R}^n$ can be split into $n+1$ parts, each having diameter strictly smaller than X. This was resolved in the negative by Kahn and Kalai [KK93]. Their spectacular combinatorial proof made Borsuk's conjecture quite popular in recent years ([Nil94] is a two-page exposition and the proof has been reproduced in several books, such as [AZ00]).

Exercises

1. Show that the antipodality assumption in Theorem 2.1.1(2.1) can be replaced by "$f(-\boldsymbol{x}) \neq f(\boldsymbol{x})$ for all $\boldsymbol{x} \in S^n$."

2. Show that the statements (2.2) and (2.3) of the Borsuk–Ulam theorem 2.1.1 are equivalent.

3. (a) Derive (1.3) from (2.2).

   (b) Derive (2.2) from (1.3).

4. Describe a surjective nullhomotopic map $S^n \to S^2$ (at least for $n = 1$ and $n = 2$).

5. (Borsuk graph) For a positive real number $\alpha < 2$, let $B(n+1, \alpha)$ be the (infinite) *Borsuk graph* with S^n as the vertex set and with two points connected by an edge iff their distance is at least α. Prove that the Borsuk–Ulam theorem is equivalent to the following statement: *For every* $\alpha < 2$, *we have* $\chi(B(n+1, \alpha)) \geq n+2$ (here χ denotes the usual chromatic number).

6. Prove that the following generalization of the Borsuk–Ulam theorem is *false* (even though it appears in the literature, according to Bourgin [Bou63, p. 337]): Whenever S^n is covered by n closed connected sets, one of them must contain a nonempty closed connected subset that is symmetric (with respect to the antipodal map).

7. Let the torus be represented as $T = S^1 \times S^1$.

   (a) Show that an analogue of the Borsuk–Ulam theorem (1.1) for maps $T \to \mathbb{R}^2$ (formulate it!) is false.

   (b) Show that it works for maps $T \to \mathbb{R}^1$.

8. Prove that the validity of (any of) the statements in the Borsuk–Ulam theorem 2.1.1 for n implies the validity of all the statements for $n-1$.

2.2 A geometric proof

We prove the version (1.2) of the Borsuk–Ulam theorem. Let $f: S^n \to \mathbb{R}^n$ be a continuous antipodal map. We want to prove that it has a zero. First we explain the idea of the proof, assuming that f is "sufficiently generic," without making the meaning of this quite precise. Then we supply a rigorous argument, involving a suitable perturbation of f.

The intuition. Let $g: S^n \to \mathbb{R}^n$ denote the "north–south projection" map; if $S^n = \{\boldsymbol{x} \in \mathbb{R}^{n+1} : x_1^2 + \cdots + x_{n+1}^2 = 1\}$, then g is given by $g(\boldsymbol{x}) = (x_1, x_2, \ldots, x_n)$. This g has exactly two zeros, namely the north pole and the south pole: $\boldsymbol{n} = (0, 0, \ldots, 0, 1)$, $\boldsymbol{s} = (0, 0, \ldots, 0, -1)$. (The important feature of g is that, obviously, it has a finite number of zeros; more precisely, it has twice an odd number of zeros.)

We consider the $(n+1)$-dimensional space $X = S^n \times [0,1]$ (a "cylinder") and the mapping $F: X \to \mathbb{R}^n$ given by $F(\boldsymbol{x}, t) = (1-t)g(\boldsymbol{x}) + tf(\boldsymbol{x})$. Geometrically, we take two copies of S^n (we can think of them as placed in $\mathbb{R}^{n+2}$), one of them with the mapping g and the other one with f. We connect the corresponding points of these two spheres by segments, and the mapping F is defined on each segment by linear interpolation. For $n = 1$, we get a cylinder as in the picture:

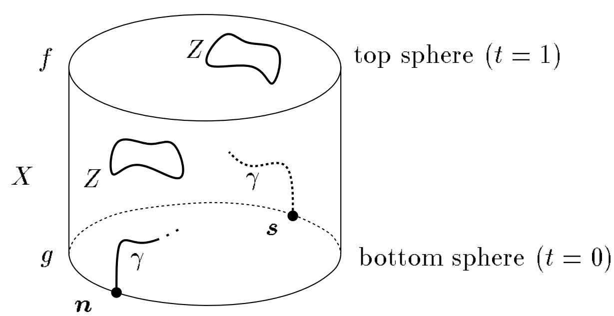

The antipodality $\boldsymbol{x} \mapsto -\boldsymbol{x}$ on S^n is extended to the map ν on X by $\nu: (\boldsymbol{x}, t) \mapsto (-\boldsymbol{x}, t)$ (note that t is unchanged). We will call ν the *antipodality on* X.

For contradiction, let us suppose that f has no zeros. We investigate the zero set $Z = F^{-1}(\mathbf{0})$. If f is sufficiently generic, then Z is a one-dimensional compact manifold, and therefore its components are cycles and paths (this is the part to be made precise later). Moreover, the endpoints of the paths lie on the bottom or top S^n ($t = 0$ or $t = 1$) and are zeros of f or g, while the cycles do not reach into the top and bottom spheres. Assuming that f has no zeros and knowing that g has only the two zeros at the poles, the only possibility is that there is a single path γ connecting $\boldsymbol{n}$ to $\boldsymbol{s}$. But, at the same time, the set Z is invariant under ν. If we follow γ from $\boldsymbol{n}$ on, the other part starting from $\boldsymbol{s}$ must behave symmetrically. But then it is easy to see that the two ends cannot meet: a symmetric path from $\boldsymbol{n}$ to $\boldsymbol{s}$ does not exist in X. We have reached a contradiction.

Note that the argument actually shows that the number of zeros of a "generic" antipodal map is twice an odd number. Indeed, the zeros of f on the top sphere are paired up by paths in Z, except for two that are connected to the zeros of g on the bottom sphere.

The real thing. A rigorous proof follows the same ideas but uses a suitable small perturbation of f. Recall that the ℓ_1-norm of a point $\boldsymbol{x} \in \mathbb{R}^n$ is $\|\boldsymbol{x}\|_1 = \sum_{i=1}^n |x_i|$. Let $\hat{S}^n = \{\boldsymbol{x} \in \mathbb{R}^{n+1} : \|\boldsymbol{x}\|_1 = 1\}$ denote the unit sphere of the ℓ_1-norm. This is the boundary of a cross-polytope (Definition 1.4.1); for example, $\hat{S}^2$ is the surface of a regular octahedron. This $\hat{S}^n$ is homeomorphic to S^n and we will consider $\hat{S}^n$ instead of S^n in the rest of the proof. The space $X = \hat{S}^n \times [0,1]$ is a union of finitely many convex polytopes (simplicial prisms). Let us call $\hat{S}^n \times \{0\}$ the *bottom sphere* and $\hat{S}^n \times \{1\}$ the *top sphere* in X.

We choose a sufficiently fine finite triangulation T of X (just how fine will be specified later) that respects the symmetry of X given by ν, in the following sense: each simplex $\sigma \in \mathsf{T}$ is mapped bijectively onto the "opposite" simplex $\nu(\sigma) \in \mathsf{T}$, and $\sigma \cap \nu(\sigma) = \emptyset$. Moreover, the triangulation T contains triangulations T_t and T_b of the top and bottom spheres, respectively, as subcomplexes, and T_t and T_b each refine the natural triangulation of $\hat{S}^n$.

We let the mapping g be an orthogonal projection of $\hat{S}^n$ into $\mathbb{R}^n$, but not in a coordinate direction, but rather in a "generic" direction, such that the two zeros $\boldsymbol{n}$ and $\boldsymbol{s}$ of g lie in the interior of n-dimensional simplices of the triangulation T_b, as is indicated in the drawing (where $n = 2$):

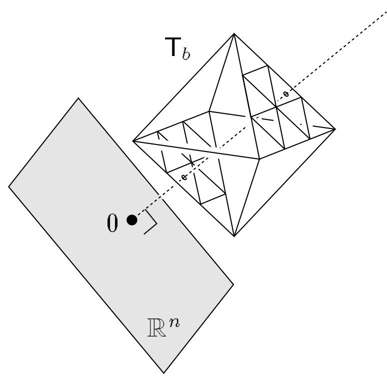

We again suppose that $f: \hat{S}^n \to \mathbb{R}^n$ has no zeros. By compactness, there is an $\varepsilon > 0$ such that $\|f(\boldsymbol{x})\| \geq \varepsilon$ for all $\boldsymbol{x} \in \hat{S}^n$. As in the informal outline, let $F(\boldsymbol{x}, t) = tg(\boldsymbol{x}) + (1-t)f(\boldsymbol{x})$, let T be a fine triangulation of X as above, and let $\bar{F}: X \to \mathbb{R}^n$ be the map that agrees with F on the vertex set $V(\mathsf{T})$ of T and is affine on each simplex of T (similar to Definition 1.5.3 of the affine extension of a simplicial map). Since F is uniformly continuous, we can assume that $\|F(\boldsymbol{y}) - \bar{F}(\boldsymbol{y})\| \leq \frac{\varepsilon}{2}$ for all $\boldsymbol{y} \in X$, provided that T is sufficiently fine. Thus,

$$\bar{F} \text{ has no zeros on the top sphere.} \tag{2.1}$$

Since our g is already affine, $\bar{F}$ coincides with g on the bottom sphere and we have

$$\begin{array}{l}\bar{F} \text{ has exactly two zeros on the bottom} \\ \text{sphere, lying in the interiors of } n\text{-dimensional} \\ \text{(antipodal) simplices of } \mathsf{T}_b.\end{array} \tag{2.2}$$

Further, let $\tilde{F}$ be a mapping arising by a sufficiently small antipodal perturbation of $\bar{F}$. Namely, we choose a suitable antipodal perturbation map $P_0: V(\mathsf{T}) \to \mathbb{R}^n$ satisfying $P_0(\nu(\boldsymbol{v})) = -P_0(\boldsymbol{v})$ for each $\boldsymbol{v} \in V(\mathsf{T})$. Further properties required of P_0 will be specified later. We extend P_0 affinely on each simplex of T, obtaining a map $P: X \to \mathbb{R}^n$, and we set $\tilde{F} = \bar{F} + P$. We note that if all values of P_0 lie sufficiently close to $\mathbf{0}$, then the perturbed map $\tilde{F}$ still has the two properties (2.1) and (2.2).

Let σ be an $(n+1)$-dimensional simplex and h an affine map $\sigma \to \mathbb{R}^n$. We say that h is *generic* if $h^{-1}(\mathbf{0})$ intersects no face of σ of dimension smaller than n. In such case, $h^{-1}(\mathbf{0})$ is either empty, or it is a segment lying in the interior of σ, with endpoints lying in the interior of two (distinct) n-faces of σ:

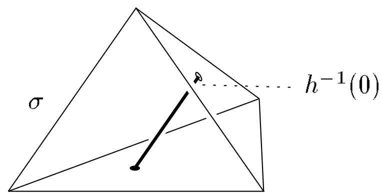

If we represent an affine map $h: \sigma \to \mathbb{R}^n$ by the $(n+2)$-tuple of values at the vertices of σ, all such maps constitute a real vector space of dimension $n(n+2)$. One can check that the set of mappings that are *not* generic is contained in a proper algebraic subvariety of this space, and so in particular, has measure zero by Sard's theorem. (Alternatively, one can check that this set is nowhere dense and use this in the sequel; see Exercise 1.)

Call a perturbed mapping $\tilde{F}: X \to \mathbb{R}^n$ *generic* if it is generic on each full-dimensional simplex of T. If T has $2N$ vertices, then the space of all possible antipodal perturbation maps P_0 on $V(\mathsf{T})$ has dimension nN (the value can be chosen freely on a set of N vertices containing no two antipodal vertices). The mappings P_0 leading to $\tilde{F}$'s that are not generic on a particular full-dimensional simplex $\sigma \in \mathsf{T}$ have measure zero in this space (here we need that $\boldsymbol{v}$ and $\nu(\boldsymbol{v})$ never lie in the same simplex of T). Therefore, arbitrarily small perturbations P_0 exist such that $\tilde{F}$ is generic.

Assuming that $\tilde{F}$ is generic and that its zeros satisfy (2.1) and (2.2), it follows that $\tilde{F}^{-1}(\mathbf{0})$ is a locally polygonal path (consisting of segments, with no branchings). This is because each n-simplex $\tau \in \mathsf{T}$ is a face of exactly two $(n+1)$-simplices $\sigma, \sigma' \in \mathsf{T}$, unless $\tau \in \mathsf{T}_t \cup \mathsf{T}_b$, in which case it is a face of exactly one $(n+1)$-simplex $\sigma \in \mathsf{T}$. Hence the components of $\tilde{F}^{-1}(\mathbf{0})$ are zero or more closed polygonal cycles (which do not intersect the top or bottom spheres) and a polygonal path γ. This γ consists of finitely many segments and it connects $\tilde{\boldsymbol{n}}$ to $\tilde{\boldsymbol{s}}$ (these are the zeros of $\tilde{F}$ on the bottom sphere).

Choose the unit of length so that γ has length 1, and let $\gamma(z)$ denote the point of γ at distance z from $\tilde{\boldsymbol{n}}$ (measured along γ; $z \in [0,1]$). Since γ is symmetric under ν, we have $\nu(\gamma(z)) = \gamma(1-z)$, and in particular, $\nu(\gamma(\frac{1}{2})) = \gamma(\frac{1}{2})$. This is impossible since the antipodality ν has no fixed points. The Borsuk–Ulam theorem is proved.

Notes. We have learned this proof from Imre Bárány, who published it, in a slightly different form, in [Bár80]. Steinlein [Ste85] gives five references for proofs of this type, all of them published between the years 1979 and 1981.

Exercises

1. (a) Let $p(x_1, x_2, \ldots, x_n) = p(\boldsymbol{x})$ be a nonzero polynomial in n variables. Show that the zero set $Z(p) = \{\boldsymbol{x} \in \mathbb{R}^n : p(\boldsymbol{x}) = 0\}$ is *nowhere dense*, meaning that any open ball B contains an open ball B' with $B' \cap Z(p) = \varnothing$.

   (b) Check that a finite union of nowhere dense sets is nowhere dense.

   (c) Let σ be an $(n+1)$-dimensional simplex; w.l.o.g. $\sigma = \operatorname{conv}\{\boldsymbol{0}, \boldsymbol{e}_1, \ldots \boldsymbol{e}_{n+1}\}$, where the $\boldsymbol{e}_i$ are the vectors of the orthonormal basis in $\mathbb{R}^{n+1}$. Let $h\colon \sigma \to \mathbb{R}^n$ be an affine map (i.e. a map of the form $\boldsymbol{x} \mapsto A\boldsymbol{x} + \boldsymbol{b}$, where A is an $n \times (n+1)$ matrix and $\boldsymbol{b} \in \mathbb{R}^n$). If each h is represented by $(h(\boldsymbol{0}), h(\boldsymbol{e}_1), \ldots, h(\boldsymbol{e}_{n+1})) \in \mathbb{R}^{(n+2)n}$, show that the maps that are not generic in the sense defined in the text above form a nowhere dense set. Hint: for each possible "cause" of non-genericity, write down a determinant that becomes 0 for all maps that are non-generic for that cause.

2.3 A combinatorial proof

Here we prove the Borsuk–Ulam theorem by a simple reduction to a purely combinatorial statement (resembling Sperner's lemma). We will be proving version (1.3), namely that a map $f\colon B^n \to \mathbb{R}^n$ that is antipodal on the boundary of B^n has a zero.

Similar to the previous section, we will replace the Euclidean ball B^n by an n-dimensional polytope. This time we take the ℓ_∞-norm unit ball $\bar{B}^n = \{\boldsymbol{x} \in \mathbb{R}^{n+1} : \|\boldsymbol{x}\|_\infty \leq 1\}$, which is the regular n-dimensional cube $\bar{B}^n = [-1, +1]^n$. Call a simplicial complex T a *special triangulation* of $\bar{B}^n$ if

- T triangulates the n-cube, $\|\mathsf{T}\| = \bar{B}^n$,
- T is a refinement of the subdivision of $\bar{B}^n$ into 2^n unit cubes by the n coordinate hyperplanes (that is, each simplex of T is fully contained in one of the 2^n orthants), and
- T is antipodally symmetric with respect to the origin: we have $\sigma \in \mathsf{T}$ iff $-\sigma \in \mathsf{T}$.

2.3.1 Lemma (Tucker's lemma). *Let the vertices of an arbitrary special triangulation T be denoted by labels $\lambda(\boldsymbol{u}) \in \{\pm 1, \pm 2, \ldots, \pm n\}$ in such a way that for the vertices $\boldsymbol{u} \in \partial \bar{B}^n$ (on the boundary) the labeling satisfies $\lambda(-\boldsymbol{u}) = -\lambda(\boldsymbol{u})$. Then there exists a 1-simplex (an edge) in T that is* **complementary**, *i.e. its two vertices are labeled by opposite numbers.*

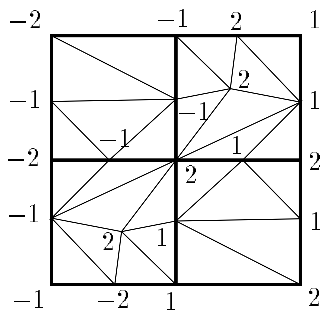

Proof of the Borsuk–Ulam theorem (1.3) from Tucker's Lemma. Assume that $f\colon \bar{B}^n \to \mathbb{R}^n$ is a map that is antipodal on the boundary and satisfies $f(\boldsymbol{x}) \neq \mathbf{0}$ everywhere. Then, from the compactness of the ball, there exists $\varepsilon > 0$ such that $\|f(\boldsymbol{x})\|_\infty \geq \varepsilon$ for all $\boldsymbol{x}$. Further, a continuous function on a compact set is uniformly continuous, and thus there exists a number $\delta > 0$ such that if the distance of some two points $\boldsymbol{x}, \boldsymbol{x}'$ does not exceed δ, then $\|f(\boldsymbol{x}) - f(\boldsymbol{x}')\|_\infty < \varepsilon$.

Let us choose a special triangulation T such that the diameter of each its simplices is at most δ. We define a labeling of the vertices of T. For $\boldsymbol{x} \in V(\mathsf{T})$, we let $i(\boldsymbol{x}) := \min\{i : |f(\boldsymbol{x})_i| \geq \varepsilon\}$, and we set

$$\lambda(\boldsymbol{x}) := \operatorname{sign}(f(\boldsymbol{x})_{i(\boldsymbol{x})}) \cdot i(\boldsymbol{x}).$$

Clearly, we have $\lambda(-\boldsymbol{x}) = -\lambda(\boldsymbol{x})$. So Tucker's lemma applies and yields a complementary edge $\boldsymbol{v}\boldsymbol{v}'$. Let $\lambda(\boldsymbol{v}) = -\lambda(\boldsymbol{v}') = i$; then $f(\boldsymbol{v})_i \geq \varepsilon$ and $f(\boldsymbol{v}')_i \leq -\varepsilon$, and hence $\|f(\boldsymbol{v}) - f(\boldsymbol{v}')\|_\infty \geq 2\varepsilon$—a contradiction.

Proof of Tucker's Lemma. Let T be a special triangulation of $\bar{B}^n$. For a simplex $\sigma \in \mathsf{T}$ we set $\operatorname{sign}(\sigma) = (\operatorname{sign}(x_1), \operatorname{sign}(x_2), \ldots, \operatorname{sign}(x_n)) \in \{+1, 0, -1\}^n$, where $\boldsymbol{x}$ is an arbitrary point of the relative interior of σ. This definition always makes sense, since a special triangulation refines orthants of $\mathbb{R}^n$ and therefore the signs of the coordinates do not change on the relative interior of σ. Let us imagine that a simplex likes to have labels corresponding to its nonzero signs: call a non-empty simplex σ *happy* if the following holds for each $i = 1, 2, \ldots, n$: if $(\operatorname{sign}(\sigma))_i = 1$, then at least one of the vertices of σ is labeled by the number i, and if $(\operatorname{sign}(\sigma))_i = -1$, then some vertex of σ is labeled by $-i$. The happy simplices are emphasized in the following picture:

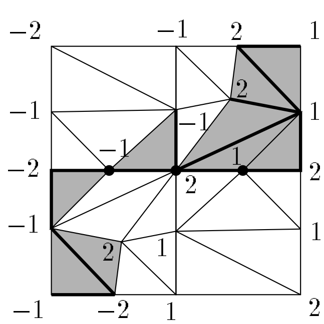

If σ has exactly k non-zero signs, that is its relative interior points have exactly k non-zero coordinates, then σ is contained in the linear span of k unit vectors, and thus its dimension is at most k. At the same time, to be happy σ needs at least k vertices with distinct labels, so the dimension of σ is at least $k-1$. Thus the dimension of a happy simplex σ for which $\operatorname{sign}(\sigma)$ has k non-zero components must be either k or $k-1$.

We define a graph G whose vertices are all happy simplices, and in which vertices $\sigma, \tau \in \mathsf{T}$ are connected by an edge if

(a) $\sigma, \tau \subset \partial \bar{B}^n = \bar{S}^{n-1}$ and $\sigma = -\tau$, or

(b) σ is a k-simplex and τ is a $(k-1)$-face of σ, such that the labels of the vertices of τ alone already make σ happy.

The simplex $\{\mathbf{0}\}$ has degree 1 in G, since it is connected exactly to the edge of the triangulation that is made happy by $\lambda(0)$. Further we prove that any other vertex σ of the graph G has degree 2 *except* when σ contains a complementary edge. Since a graph cannot contain only one vertex of odd degree, this will establish Tucker's lemma. (Verify in our sketch, or even better in your own example, that the graph contains a path that connects $\mathbf{0}$ to some happy simplex that contains a complementary edge!)

Let $\operatorname{sign}(\sigma)$ have k nonzero components, so $\dim \sigma$ is k or $k-1$. We distinguish these two cases.

1. Suppose that σ is a $(k-1)$-simplex. Here we have two subcases:

   1.1 σ does not lie on the boundary of $\bar{B}^n$. Then we claim that it is a face of exactly two k-simplices that are made happy by the k obligatory labels of σ. Indeed, any k-simplex made happy by the labels of σ must be contained in the k-dimensional coordinate subspace $L_\sigma := \{\boldsymbol{x} \in \mathbb{R}^n : x_i = 0 \text{ for all } i \text{ with } \operatorname{sign}(\sigma)_i = 0\}$. The intersection $L_\sigma \cap \bar{B}^n$ is a k-cube and the simplices of T contained in L_σ triangulate it. Now if σ is a non-boundary $(k-1)$-dimensional simplex in a triangulation of a k-cube, it is adjacent to precisely two k-simplices.

   1.2 σ is on the boundary of $\bar{B}^n$. Then it has $-\sigma$ as one neighbor. Arguing similar to the previous case, we get that σ is a face of exactly one k-simplex made happy by the labels of σ, and this is the other neighbor.

2. If σ is a k-simplex, it has k obligatory labels and one extra label. Note that in this case σ cannot lie on the boundary. The possible cases are

   2.1 The extra label repeats one of the obligatory labels. Then σ is adjacent to two of its $(k-1)$-faces.

   2.2 The extra label is the negative of some of the obligatory labels, but then we have a complementary edge.

   2.3 The extra label is a number i such that $\pm i$ does not occur among the obligatory labels. Then one of the neighbors of σ is the unique $(k-1)$-face with the obligatory labels. Moreover, σ is a face of exactly one

$(k+1)$-dimensional simplex σ' made happy by the labels of σ. We enter that σ' if we go from an interior point of σ in the direction of the $x_{|i|}$-axis, in the positive direction for $i > 0$ and in the negative direction for $i < 0$. Thus, $\mathrm{sign}(\sigma')$ coincides with $\mathrm{sign}(\sigma)$ everywhere except position $|i|$, where σ has 0 and σ' has $\mathrm{sign}(i)$.

So for each possibility we have exactly two neighbors, which yields a contradiction.

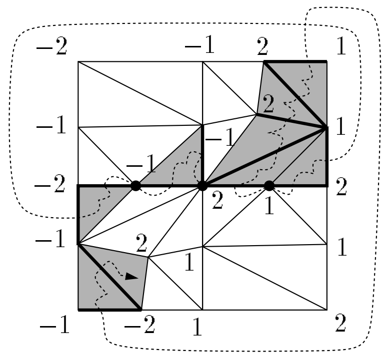

Thus the graph that we defined leads us from the vertex in the origin to a simplex with a complementary edge.

Remark. Why is version (1.3) of the Borsuk–Ulam theorem especially suitable for a parity-based argument as above? This is because a generic mapping f as in (1.3) has an *odd* number of zeros (while in the "basic" version with an antipodal map $S^n \to \mathbb{R}^n$, the zeros come in pairs).

Notes. Tucker's lemma is from [Tuc46] (this paper contains a 2-dimensional version, and a version for arbitrary dimension appeared in the book [Lef49]). The presented proof follows Freund & Todd [FT81]. Steinlein's survey [Ste85] lists over 10 other references with combinatorial proofs based on similar ideas.

3

Direct Applications of Borsuk–Ulam

3.1 The ham sandwich theorem

This is a well-known geometric statement with many interesting consequences. The informal statement that gave the ham sandwich theorem its name is: *For every sandwich made of ham, cheese, and bread, there is a straight cut that simultaneously halves the ham, the cheese, and the bread.* The mathematical ham sandwich theorem says that any d distributions of mass in $\mathbb{R}^d$ can be simultaneously bisected by a hyperplane:

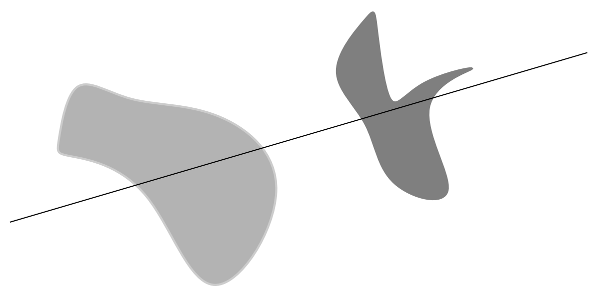

First we prove a statement about equipartitioning suitable finite Borel measures $\mu_1, \ldots, \mu_d$ in $\mathbb{R}^d$. A finite Borel measure μ on $\mathbb{R}^d$ is a measure on $\mathbb{R}^d$ such that all open subsets of $\mathbb{R}^d$ are measurable and $\mu(\mathbb{R}^d) < \infty$. An example the reader may want to think of is a measure given as the restriction of the usual Lebesgue measure to a compact subset of $\mathbb{R}^d$. That is, $A \subset \mathbb{R}^d$ is compact with $\lambda^d(A) > 0$, where λ^d denotes the d-dimensional Lebesgue measure, and $\mu(X) = \lambda^d(X \cap A)$ for all (Lebesgue measurable) sets $X \subseteq \mathbb{R}^d$.

3.1.1 Theorem (Ham sandwich theorem for measures). *Let $\mu_1, \mu_2, \ldots, \mu_d$ be finite Borel measures on $\mathbb{R}^d$ such that every hyperplane has measure 0 for each of the μ_i (in the sequel, we refer to such measures as "mass distributions"). Then there exists a hyperplane h such that*

$$\mu_i(h^+) = \tfrac{1}{2}\,\mu_i(\mathbb{R}^d) \quad \text{for } i = 1, 2, \ldots, d,$$

where h^+ denotes one of the halfspaces defined by h.

Proof. Let $\boldsymbol{u} = (u_0, u_1, \ldots, u_d)$ be a point of the sphere S^d. If at least one of the components $u_1, u_2, \ldots, u_d$ is nonzero, we assign to the point $\boldsymbol{u}$ the halfspace

$$h^+(\boldsymbol{u}) := \{(x_1, \ldots, x_d) \in \mathbb{R}^d : u_1x_1 + \cdots + u_dx_d \le u_0\}.$$

Obviously antipodal points of S^d correspond to opposite halfspaces. For a $\boldsymbol{u}$ of the form $(u_0, 0, 0, \ldots, 0)$ (where $u_0 = \pm 1$), we have by the same formula

$$\begin{aligned} h^+((1,0,\ldots,0)) &= \mathbb{R}^d \\ h^+((-1,0,\ldots,0)) &= \varnothing. \end{aligned}$$

We define a function $f\colon S^d \to \mathbb{R}^d$ by

$$f_i(\boldsymbol{u}) = \mu_i(h^+(\boldsymbol{u})).$$

It is easily checked that if we have $f(\boldsymbol{u}_0) = f(-\boldsymbol{u}_0)$ for some $\boldsymbol{u}_0 \in S^d$, then the boundary of the halfspace $h^+(\boldsymbol{u}_0)$ is the desired hyperplane (clearly it cannot happen that $f((1,0,\ldots,0)) = f((-1,0,\ldots,0))$, so $h^+(\boldsymbol{u}_0)$ is indeed a halfspace). For an application of the Borsuk–Ulam theorem it remains to show that f is continuous. This is quite intuitive but a rigorous argument is perhaps not so obvious.

Let $(\boldsymbol{u}_n)_{n=1}^{\infty}$ be a sequence of points of S^d converging to $\boldsymbol{u}$; we need to show that $\mu_i(h^+(\boldsymbol{u}_n)) \to \mu_i(h^+(\boldsymbol{u}))$. We note that if a point $\boldsymbol{x}$ is not on the boundary of $h^+(\boldsymbol{u})$, then for all sufficiently large n, we have $\boldsymbol{x} \in h^+(\boldsymbol{u}_n)$ if and only if $\boldsymbol{x} \in h^+(\boldsymbol{u})$. So if f denotes the characteristic function of $h^+(\boldsymbol{u})$ ($f(\boldsymbol{x}) = 1$ for $x \in h^+(\boldsymbol{u})$ and $f(\boldsymbol{x}) = 0$ for $x \notin h^+(\boldsymbol{u})$) and f_n is the characteristic function of $h^+(\boldsymbol{u}_n)$, we have $f_n(\boldsymbol{x}) \to f(\boldsymbol{x})$ for all $\boldsymbol{x} \notin \partial h^+(\boldsymbol{u})$. Since $\partial h^+(\boldsymbol{u})$ has μ_i-measure 0 by the assumption, the f_n converge to f μ_i-almost everywhere. By Lebesgue's dominated convergence theorem (see e.g. Rudin [Rud74, Theorem 1.34]), we thus have $\mu_i(h^+(\boldsymbol{u}_n)) = \int f_n \boldsymbol{d}\mu_i \to \int f \boldsymbol{d}\mu_i = \mu_i(h^+(\boldsymbol{u}))$, as all the f_n are dominated by the constant 1, which is integrable since μ_i is finite. (It is not difficult to prove the particular case of the dominated convergence theorem needed here directly.)

Sometimes we need to partition masses concentrated at finitely many points. Then the following version of the ham sandwich theorem can be useful:

3.1.2 Theorem (Ham sandwich theorem for point sets).
Let $A_1, A_2, \ldots, A_d \subset \mathbb{R}^d$ be finite point sets. Then there exists a hyperplane h that bisects each A_i.

Here "h bisects A_i" means that both the open halfspaces defined by h contain at most $\frac{1}{2}|A_i|$ points of A_i.

Proof from Theorem 3.1.1. First suppose that each A_i has odd cardinality and that $A_1 \dot\cup A_2 \dot\cup \ldots \dot\cup A_d$ is in general position, meaning that no two points of different A_i coincide and no $d{+}1$ points lie on a common hyperplane. Let A_i^ε arise from A_i by replacing each point by a solid ball of radius ε centered at that point, and choose $\varepsilon > 0$ so small that no $d{+}1$ balls of $\bigcup A_i^\varepsilon$ can be intersected by a common hyperplane. Let h be a hyperplane simultaneously bisecting the sets A_i^ε. Since A_i^ε has an odd number of balls, h must intersect at least one of them, and since at most d balls are intersected altogether, h intersects exactly one ball of A_i^ε. Moreover, this ball is split in half by h, and so h passes through its center. Thus h bisects each A_i.

Next, let the A_i still have odd cardinality but their position can be arbitrary. We use a perturbation argument. For every $\eta > 0$, let $A_{i,\eta}$ arise from A_i by moving each point by at most η in such a way that $\bigcup_{i=1}^{d} A_{i,\eta}$ is in general position. Let h_η bisect the $A_{i,\eta}$. If we write $h_\eta = \{\boldsymbol{x} \in \mathbb{R}^d : \langle \boldsymbol{a}_\eta, \boldsymbol{x} \rangle = b_\eta\}$, where $\boldsymbol{a}_\eta$ is a unit vector, then the b_η lie in a bounded interval, and so by compactness, there exists a cluster point $(\boldsymbol{a}, b) \in \mathbb{R}^{d+1}$ of the pairs $(\boldsymbol{a}_\eta, b_\eta)$ as $\eta \to 0$. Let h be the hyperplane determined by the equation $\langle \boldsymbol{a}, \boldsymbol{x} \rangle = b$. Let us consider a sequence $\eta_1 > \eta_2 > \cdots$ converging to 0 such that $(\boldsymbol{a}_{\eta_j}, b_{\eta_j}) \to (\boldsymbol{a}, b)$. If a point $\boldsymbol{x}$ lies at distance $\delta > 0$ from h, then it also lies at distance at least $\frac{1}{2}\delta$, say, from h_{η_j} for all sufficiently large j. Therefore, if there are k points of A_i in one of the open halfspaces determined by h then, for all j large enough, the corresponding open halfspace determined by h_{η_j} contains at least k points of A_{i,η_j}. It follows that h bisects all the A_i.

Finally, if some of the A_i has an even number of points, we delete one arbitrarily chosen point from each even-size A_i and bisect the resulting odd-size sets. Adding the deleted points back cannot spoil the bisection, as is easy to check from the definition of bisection.

3.1.3 Corollary (Ham sandwich theorem, general position version). *Let $A_1, A_2, \ldots, A_d \subset \mathbb{R}^d$ be disjoint finite point sets in general position (such that no more than d points of $A_1 \dot\cup \cdots \dot\cup A_d$ are contained in any hyperplane).*

Then there exists a hyperplane h that bisects each A_i, such that there are exactly $\lfloor \frac{1}{2} |A_i| \rfloor$ points from A_i in each of the open halfspaces defined by h, and at most one point of A_i on the hyperplane h (which happens if A_i has odd cardinality).

Proof. We start with an arbitrary ham sandwich cut hyperplane h according to Theorem 3.1.2. We fix the coordinate system so that h is the horizontal hyperplane $x_d = 0$. Let $B := h \cap (A_1 \cup \cdots \cup A_d)$; B consists of at most d affinely independent points. We want to move h slightly so that it is as in the corollary (i.e. only one point of each odd-size A_i stays on it). Since the points of B are affinely independent we can make each of them stay on h or go below or above it, whatever we decide.

To see this, we add $d - |B|$ new points to B so that we obtain a d-point affinely independent $C \subset h$. For each $\boldsymbol{a} \in C$, we choose a point $\boldsymbol{a}'$: either $\boldsymbol{a}' = \boldsymbol{a}$ (for the new points $\boldsymbol{a}$ and for those points of B that should stay on h), or $\boldsymbol{a}' = \boldsymbol{a} + \varepsilon \boldsymbol{e}_d$, or $\boldsymbol{a}' = \boldsymbol{a} - \varepsilon \boldsymbol{e}_d$. We let $h' = h'(\varepsilon)$ be the hyperplane determined by the d points $\boldsymbol{a}'$, $\boldsymbol{a} \in C$. For all sufficiently small $\varepsilon > 0$, the $\boldsymbol{a}'$ remain affinely independent (so that $h'(\varepsilon)$ is well-defined) and the motion of $h'(\varepsilon)$ is continuous in ε. We can thus guarantee that for all sufficiently small $\varepsilon > 0$, h' is as required in the corollary.

Equipartition theorems. Using the 2-dimensional ham sandwich theorem, it is easy to show that any mass distribution in the plane can be dissected into 4 equal parts by 2 lines:

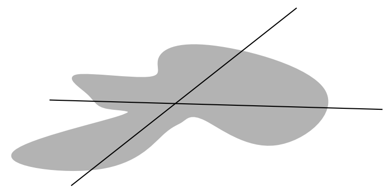

As a natural generalization, one can ask whether any mass distribution in $\mathbb{R}^3$ can be partitioned into $2^3 = 8$ equal pieces by 3 planes, or, more generally, if any mass distribution in $\mathbb{R}^d$ can be dissected into 2^d pieces of equal measure by d hyperplanes. For $d = 3$, this is possible (although not as simple as the planar case; see Edelsbrunner [Ede87, Sect. 4.4]). But in dimension 5 and higher, such an equipartition theorem fails: it is in general *impossible* to cut a set in $\mathbb{R}^5$ into 32 equal parts by 5 hyperplanes. For this, note that any hyperplane cuts the moment curve in $\mathbb{R}^5$ in at most 5 distinct points; hence any set of 5 hyperplanes cuts the moment curve in at most 25 distinct points, subdividing it into at most 26 parts. So if we take a piece of the moment curve, it is disjoint with at least 6 of the 32 open orthants determined by 5 hyperplanes, and hence it cannot be equipartitioned. This example uses a one-dimensional measure along the moment curve; an example obtained by restricting the Lebesgue measure to suitable small balls requires a little more work (Avis [Avi85]; also see Edelsbrunner [Ede87, Sect. 4.6].) It is not known whether a dissection into 16 parts of the same size by 4 hyperplanes is possible in $\mathbb{R}^4$, and it is a challenging open problem where many of the "usual" topological approaches seem to fail.

There are numerous results on equipartitions of measures; some of them will be mentioned in the remarks below and in the exercises.

Notes. According to [Ste85], the ham sandwich theorem was conjectured by Steinhaus and proved by Banach.

The ham sandwich theorem in $\mathbb{R}^d$ can be, and often is, proved from the $(d-1)$-dimensional Borsuk–Ulam theorem. For every direction $\boldsymbol{u} \in S^{d-1}$, one chooses the hyperplane $h(\boldsymbol{u})$ perpendicular to $\boldsymbol{u}$ that bisects the dth measure, and defines the function to $\mathbb{R}^{d-1}$ as the parts of the 1st through $(d-1)$st measures contained in $h(\boldsymbol{u})^+$. But, to guarantee uniqueness of $h(\boldsymbol{u})$, and thus derive the existence of a continuous function h, one needs a stronger assumption on the measures.

Dol'nikov [Dol'92] and, independently, Živaljević & Vrećica [ŽV90] proved, by more advanced topological means, a nice generalization of the ham sandwich theorem. For any $k+1$ mass distributions in $\mathbb{R}^d$ there exists a k-flat f (i.e. a k-dimensional affine subspace of $\mathbb{R}^d$) such that any hyperplane passing through f has at least $\frac{1}{d-k+1}$ of the i-th mass on each side, for all $i = 1, 2, \ldots, k+1$. The ham sandwich theorem is obtained for $k = d-1$. The case $k = 0$ is another classical result known as the *centerpoint theorem* (see e.g. [Ede87]).

Mass partition theorems. Results on partitioning of one or several masses in $\mathbb{R}^d$ into prescribed parts by given geometric objects are almost always proved by topological methods. Interest in such results was stimulated by applications in computer science, for example in the so-called *geometric range searching*; see [Mat95] [AE98]. (In this area, though, approximate partitioning is usually sufficient, and the classical mass partitioning results were eventually superseded by random sampling and related methods.)

Concerning the problem of dissecting a measure in $\mathbb{R}^4$ into 16 equal parts by 4 hyperplanes, we remark that partitioning of 16 points placed on the moment curve is always possible. This is equivalent to the existence of a uniform Gray code in the 4-dimensional cube: there is a Hamiltonian circuit in the graph of the 4-cube that uses the same number of edges (4) from each parallel class. In fact, Robinson and Cohen [RC81] showed that a uniform Gray code in C_n exists if and only if n is a power of 2. Ramos [Ram96] gives several new results on partitioning of m mass distributions in $\mathbb{R}^d$ into 2^k equal pieces by k hyperplanes. Also see the survey by Živaljević [Živ98] for a description of still newer results of Petrović et al. in this direction, obtained using obstruction theory.

Recently, several results have been proved concerning partitions by *k-fans*, i.e. by k semilines emanating from a common point in the plane (the point may also be at infinity, i.e. we may have k parallel lines; in this case, both the unbounded parts of the plane together form one sector). Answering a question of Kaneko and Kano [KK99], several authors [IUY00] [Sak] [BKS00] have shown that two mass distributions in the plane can be simultaneously equipartitioned by a 3-fan, in such a way that the resulting 3 sectors are convex. For example, a planar convex body can be cut by a 3-fan so that both the area and the perimeter are divided equitably (this special "cake cutting" case was shown in [AKK$^+$00]):

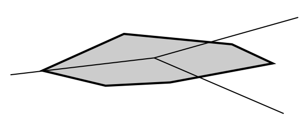

Several results on partitions of m measures by k-fans are proved in [BM01], including some cases where the partition is not into equal parts; for example, any 2 measures can be simultaneously partitioned in ratio $1 : 1 : 1 : 2$ by a 4-fan (without any convexity requirements). Later the possibility of equipartition of 2 measures by a 4-fan was shown as well [BM], but challenging problems remain open; for instance, can any 2 measures be partitioned by a 4-fan in any prescribed ratio? Further progress on this question was recently announced by Živaljević and Vrećica.

Somewhat related results (dealing with a single measure), including some higher-dimensional ones, were given earlier by Makeev [Mak88]. He established the existence of 6-partitions by suitable cones in $\mathbb{R}^3$; for example, for any mass distribution in $\mathbb{R}^3$, there is a cube C such that the six infinite cones with apex in the center of C and with the facets of C as bases form an equipartition. Živaljević and Vrećica [ŽV01] also proved several higher-dimensional results, such as that given a simplex Δ in $\mathbb{R}^d$ and a point $x \in \operatorname{int} \Delta$, any given mass distribution can be dissected into $d+1$ parts with arbitrary prescribed ratios by a suitable translation of the $d+1$ cones with apex x given by the facets of Δ.

Another interesting equipartitioning result is Schulman's [Sch93] "cobweb partition theorem": every bounded set of finite measure in $\mathbb{R}^2$ has a partition into 8 equally large parts by a cobweb as in the figure.

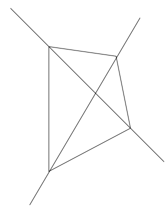

Exercises

1. Consider 3 mass distributions in the plane which, moreover, assign measure 0 to each circle. Prove that they can be simultaneously halved by a circle or by a straight line. (This is a special case of results of Stone & Tukey; see [Bre93, p. 243].)

2. Show that 1 : 1 is the only ratio such that any two compact sets in the plane can be simultaneously partitioned by a line in that ratio.

3. (a) Find 4 measures in the plane that cannot be simultaneously bisected by a 2-fan.

   (b) Find 3 measures in the plane that cannot be simultaneously equipartitioned by a 3-fan.

   (c) Find 2 measures in the plane that cannot be simultaneously equipartitioned by a 5-fan.

   See [BM01] for a detailed solution.

3.2 On multicolored partitions and necklaces

Multicolored partitions. Here is one nice and simple consequence of the (discrete) ham sandwich theorem:

3.2.1 Theorem (Akiyama & Alon [AA89]). *Consider d n-point sets $A_1, \ldots, A_d$ in general position in $\mathbb{R}^d$; imagine that the points of A_1 are red, the points of A_2 blue, etc. (each A_i has its own color). Then the points of the union $A_1 \cup \cdots \cup A_d$ can be partitioned into "rainbow" d-tuples (each d-tuple contains one point of each color) with pairwise disjoint convex hulls.*

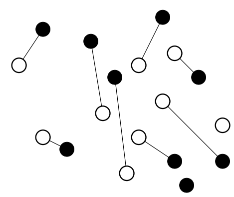

(In our drawing we didn't quite manage to find a correct pairing.)

Proof. We proceed by induction on n. If $n > 1$ is odd, there is a hyperplane h bisecting each A_i and containing exactly one point of each color. We let the points in h form one d-tuple and use induction for the subsets in the open halfspaces. For n even, we invoke the general-position version of the ham-sandwich theorem (Corollary 3.1.3), which guarantees a bisecting hyperplane that avoids all the A_i.

Remark. For $d = 2$ the theorem can be proved directly (Exercise 1). No direct (non-topological) proof is known in higher dimensions.

Division of a necklace. Two thieves have stolen a precious necklace of nearly immeasurable value, not only because of the precious stones (diamonds, saphirs, rubies, etc.), but also because these are set in pure platinum. The thieves do not know the values of the stones of various kinds, and so they want to divide the stones each kind evenly. In order to waste as little platinum as possible, they want to achieve this by as few cuts as possible (admittedly, this mathematical model of thieves is not very realistic).

We assume that the necklace is open (with two ends) and that there are d different kinds of stones, even number of each kind. It is easy to see that at least d cuts may be necessary: place the stones of the first kind first, then the stones of the second kind, and so on. The necklace theorem shows that this is the worst what can happen.

3.2.2 Theorem (Necklace theorem). *Every (open) necklace with d kinds of stones can be divided between two thieves using no more than d cuts.*

So for the necklace in our picture, 3 cuts should suffice:

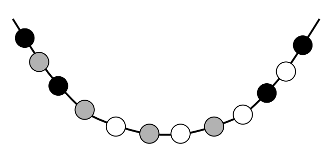

Surprisingly, all known proofs of this theorem are topological.

First proof: by ham sandwich. We place the considered necklace into $\mathbb{R}^d$ along the moment curve. Let $\gamma(t) = (t, t^2, \ldots, t^d)$ be the parametric expression of the moment curve γ. If the necklace has n stones, we define

$$A_i = \{\gamma(k) : \text{the } k\text{th stone is of the } i\text{th kind, } k = 1, 2, \ldots, n\}.$$

Let us also call the points of A_i the stones of the ith kind. By the (general position discrete) ham sandwich theorem 3.1.3, there exists a hyperplane h simultaneously bisecting each A_i. This h cuts the moment curve, and the necklace lying along it, in at most d places. All the sets A_i were assumed to be of even size, so h contains no stones, and these cuts are as required in the necklace problem.

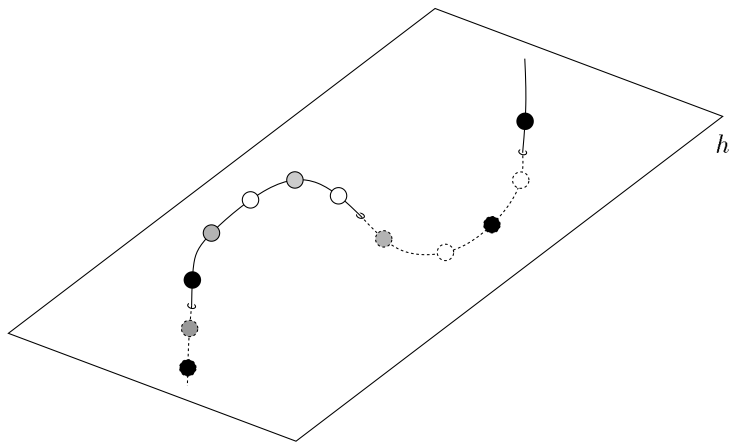

Second proof. We reproduce another proof as well, whose clever encoding of the divisions of the necklace by points of the sphere is of independent interest.

First we note that the result follows from a continuous version. By a *continuous probability measure* on $[0,1]$ we mean a probability measure μ on $[0,1]$ such that $\int_0^x \mathrm{d}\mu$ is continuous in x.

3.2.3 Theorem (Continuous necklace; Hobby–Rice theorem [HR65]). *Let $\mu_1, \mu_2, \ldots, \mu_d$ be continuous probability measures on $[0,1]$. Then there exists a partition of $[0,1]$ into $d+1$ intervals $I_0, I_1, \ldots, I_d$ (using d cut points) and signs $\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_d \in \{-1,+1\}$ with*

$$\sum_{j=0}^{d} \varepsilon_j \cdot \mu_i(I_j) = 0 \quad \text{for } i = 1, 2, \ldots, d.$$

It should be clear that it suffices to prove this result in the special case where $\varepsilon_j = (-1)^j$, since a cut point at which the sign doesn't change may be removed. However, the proof we give below does not have a natural restriction to the special case.

We also note that the Hobby–Rice theorem can be derived from the continuous ham sandwich theorem, by an argument similar to the above proof of the necklace theorem.

Proof of the necklace theorem from the continuous version. Let us have t_i stones of the i-th kind, $n := \sum_{i=1}^{d} t_i$. We imagine the necklace on the interval $[0,1]$; the k-th stone corresponds to the segment $[\frac{k-1}{n}, \frac{k}{n})$. First we define characteristic functions $f_i(x)\colon [0,1] \to \{0,1\}$ for $x \in [\frac{k-1}{n}, \frac{k}{n})$, by

$$f_i(x) = \begin{cases} 1 & \text{if the } k\text{-th stone of the necklace is of the } i\text{-th kind} \\ 0 & \text{otherwise.} \end{cases}$$

Each function f_i defines a measure μ_i on $[0,1]$, by $\mu_i(A) := \frac{n}{t_i}\int_A f_i(x)\,\mathrm{d}x$. Thus $\mu_i(A)$ denotes the fraction of stones of the i-th kind that is on the part A of the necklace.

For these μ_i, we find a division as in the continuous necklace theorem (the first thief gets the intervals with "+" signs and the second those with "−"). This division is fair but it can be nonintegral (i.e., some stones would have to be cut). We use a rounding procedure. We proceed by induction on the number of "nonintegral" cuts. If a cut subdivides a stone of the i-th type, then either the cut is unnecessary, or there is another cut through a stone of type i, and we move one cut to the right, and the other cut to the left, without changing the balance.

Proof of the continuous necklace theorem. With every point $\boldsymbol{x} = (x_1, x_2, \ldots, x_d, x_{d+1}) \in S^d$ we associate a division of the interval $[0,1]$ into $d+1$ parts, of lengths $x_1^2, x_2^2, \ldots, x_{d+1}^2$: that is, with $\boldsymbol{x}$ we associate the cuts at the points $z_i := x_1^2 + \cdots + x_i^2$, where $0 = z_0 \le z_1 \le \cdots \le z_d \le z_{d+1} = 1$. The sign ε_j for the interval $I_j = [z_j, z_{j+1}]$ is chosen as $\operatorname{sign}(x_j)$. This defines a continuous function $g\colon S^d \to \mathbb{R}^d$:

$$g_i(\boldsymbol{x}) := \sum_{j=1}^{d+1} \operatorname{sign}(x_j) \cdot \mu_i([z_{j-1}, z_j]).$$

In words, $g_i(\boldsymbol{x})$ is the amount of i-stone given to the first thief minus the amount of i-stone allocated to the second thief. This function is clearly antipodal. Thus, an $\boldsymbol{x} \in S^d$ exists with $g(\boldsymbol{x}) = 0$. This $\boldsymbol{x}$ encodes a just division.

For a solution of a similar problem with more than two thieves, the proof via the ham sandwich theorem doesn't seem to work anymore. The second proof can be generalized but the Borsuk–Ulam theorem needs to be generalized as well: instead of the sphere we have to use a different "configuration space" that admits a symmetry of higher order. The necklace problem with several thieves will be discussed in Section 6.4.

Notes. The necklace theorem was first proved by Goldberg & West [GW85]. Alon & West [AW86] found a new elegant proof, essentially the second proof given above. The proof of the necklace theorem via the ham sandwich theorem was noted by Alon (private communication) and also by Ramos [Ram96]. The continuous necklace theorem was proved by Hobby and Rice [HR65], earlier than the discrete version, and in a completely different context—but the proof is also based on the Borsuk–Ulam theorem.

Exercises

1. Prove the planar case ($d = 2$) of Theorem 3.2.1 by considering a perfect red-blue matching with the minimum possible total length of the edges.

3.3 Kneser's conjecture

One of the earliest and most spectacular applications of topological methods in combinatorics is Lovász' 1978 proof [Lov78] of the so-called Kneser conjecture. Kneser posed the following problem in 1955:

> **Aufgabe 360:** k und n seien zwei natürliche Zahlen, $k \leqq n$; N sei eine Menge mit n Elementen, N_k die Menge derjenigen Teilmengen von N, die genau k Elemente enthalten; f sei eine Abbildung von N_k auf eine Menge M, mit der Eigenschaft, daß $f(K_1) \neq f(K_2)$ ist falls der Durchschnitt $K_1 \cap K_2$ leer ist; $m(k, n, f)$ sei die Anzahl der Elemente von M und $m(k, n) = \operatorname{Min}_f m(k, n, f)$. Man beweise: Bei festem k gibt es Zahlen $m_0 = m_0(k)$ und $n_0 = n_0(k)$ derart, daß $m(k,n) = n - m_0$ ist für $n \geqq n_0$; dabei ist $m_0(k) \geqq 2k-2$ und $n_0(k) \geqq 2k-1$; in beiden Ungleichungen ist vermutlich das Gleichheitszeichen richtig.
>
> Heidelberg. MARTIN KNESER.

Let k and n be two natural numbers, $k \leq n$; let N be a set with n elements, N_k the set of all subsets of N with exactly k elements; let f be a map from N_k to a set M with the property that $f(K_1) \neq f(K_2)$ if the intersection $K_1 \cap K_2$ is empty; let $m(k, n, f)$ be the number of elements of M, and $m(k, n) = \min_f m(k, n, f)$. Prove that: for fixed k there are numbers $m_0 = m_0(k)$ and $n_0 = n_0(k)$ such that $m(k, n) = n - m_0$ for $n \geq n_0$; here $m_0(k) \geq 2k-2$ and $n_0(k) \geq 2k-1$; both inequalities probably hold with equality.

We will use slightly different notation, and recast this in a graph-theoretical language. We take $N = [n]$, we write $\binom{[n]}{k}$ instead of N_k for the collection of all k-subsets of $[n]$, we take $\binom{[n]}{k}$ as the vertex set of a graph, and we connect two vertices by an edge if the corresponding k-sets are disjoint. Then the mapping f becomes a *coloring* of the graph, where M is the set of colors, and Kneser asks for the *chromatic number* of the graph!

We recall that a *(proper) k-coloring* of a graph $G = (V, E)$ is a mapping $c\colon V \to [k]$ such that $c(u) \neq c(v)$ whenever $\{u, v\} \in E$ is an edge. The *chromatic number of* G, denoted by $\chi(G)$, is the smallest k such that G has a k-coloring.

Let X be a finite ground set and let $\mathcal{S} \subseteq 2^X$ be a set system. The *Kneser graph of* $\mathcal{S}$, denoted by $\mathrm{KG}(\mathcal{S})$, has $\mathcal{S}$ as the vertex set, and two sets $S_1, S_2 \in \mathcal{S}$ are adjacent iff $S_1 \cap S_2 = \emptyset$. In symbols,

$$\mathrm{KG}(\mathcal{S}) = \Big(\mathcal{S}, \{\{S_1, S_2\} : S_1, S_2 \in \mathcal{S},\ S_1 \cap S_2 = \emptyset\}\Big).$$

Let $\mathrm{KG}_{n,k}$ denote the Kneser graph of the system $\mathcal{S} = \binom{[n]}{k}$ (all k-element subsets of $[n]$). Then Kneser's conjecture is $\chi(\mathrm{KG}_{n,k}) = n-2k+2$ for $n \geq 2k-1$.

3.3.1 Examples.

- $\mathrm{KG}_{n,1}$ is the complete graph K_n with $\chi(K_n) = n$.
- $\mathrm{KG}_{2k-1,k}$ is a graph with no edges, and so $\chi(\mathrm{KG}_{2k-1,k}) = 1$.
- $\mathrm{KG}_{2k,k}$ is a matching (every set is adjacent only to its complement) and $\chi(\mathrm{KG}_{2k,k}) = 2$ for all $k \geq 1$.
- The first interesting example is $\mathrm{KG}_{5,2}$, which turns out to be the ubiquitous *Petersen graph*:

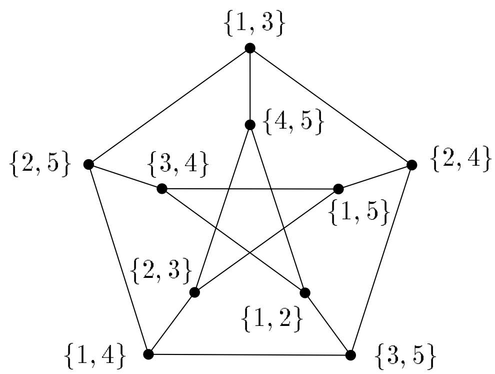

This graph serves as a "(counter)example for almost everything" in Graph Theory (see [CW85], [CHW92], [HS93] and the references given there). Check that 3 colors suffice and are necessary!

Kneser's conjecture, which is a theorem since 1978, can be restated as follows:

3.3.2 Theorem (Kneser's conjecture). *For all $k > 0$ and $n \geq 2k - 1$, the chromatic number of the Kneser graph $\mathrm{KG}_{n,k}$ is $\chi(\mathrm{KG}_{n,k}) = n - 2k + 2$.*

The Kneser graphs $\mathrm{KG}_{n,k}$ are very interesting examples of graphs with high chromatic number. For example, note that for $n = 3k-1$, they have no triangles, and yet the chromatic number is $k+1$. One of the main reasons of their importance, and also probably a reason why the proof of Kneser's conjecture is difficult, is that there is a large gap between the chromatic number and the *fractional chromatic number*. (There are *very* few examples of such graphs known.) The fractional chromatic number $\chi_f(G)$ of a graph G is defined as the infimum (actually minimum) of the fractions $\frac{a}{b}$ such that $V(G)$ can be covered by a independent sets in such a way that every vertex is covered at least b times. We always have $\chi_f(G) \leq \chi(G)$, and many methods for estimating $\chi(G)$ from below actually estimate $\chi_f(G)$—which means that they don't give good estimates for graphs that have high chromatic number $\chi(G)$, but low fractional chromatic number $\chi_f(G)$, as in the case of the Kneser graphs.

For example, the well-known lower bound in terms of the maximal size of independent sets, $\chi(G) \geq |V(G)|/\alpha(G)$, is just a part of the chain

$$\frac{|V|}{\alpha(G)} \quad \leq \quad \chi_f(G) \quad \leq \quad \chi(G),$$

where $\alpha(G)$, the *independence number* of G, is the maximum size of an independent set in G. But for the Kneser graph, we have $\chi_f(\mathrm{KG}_{n,k}) = \frac{n}{k}$(Exercise 1). So, for example, $\chi_f(\mathrm{KG}_{3k-1,k}) < 3$.

It is fairly easy to show that the chromatic number of $\mathrm{KG}_{n,k}$ cannot be larger than $n-2k+2$.

Upper bound for the chromatic number. We color the vertices of the Kneser graph by

$$\chi(F) := \min\{\min(F), n-2k+2\}.$$

This assigns a color $\chi(F) \in \{1, 2, \ldots, n-2k+2\}$ to each subset $F \in \binom{[n]}{k}$. If two sets F, F' get the same color $\chi(F) = \chi(F') = i < n-2k+2$, then they cannot be disjoint since they both contain the element i. If the two k-sets both get the color $n-2k+2$, then they are both contained in the set $\{n-2k+2, \ldots, n\}$, which has only $(2k-1)$ elements, and hence they cannot be disjoint either.

All known proofs of the tight lower bound for $\chi(\mathrm{KG}_{n,k})$ are topological or at least imitate the topological proofs. We begin with one of the simplest known proofs, found by Bárány [Bár78] soon after the announcement of Lovász' breakthrough. It is based on the following geometric lemma.

3.3.3 Lemma (Gale's lemma [Gal56]). *For every $d \geq 0$ and every $k \geq 1$, there exists a set $V \subset S^d$ of $2k+d$ points such that every open hemisphere of S^d contains at least k points of V.*

First let us see how this implies Kneser's conjecture.

First proof of Kneser's conjecture. Let us consider the Kneser graph $\mathrm{KG}_{n,k}$ and set $d := n-2k$. Let $V \subset S^d$ be the set as in Gale's lemma 3.3.3. Let us suppose that the vertex set of $\mathrm{KG}_{n,k}$ is $\binom{V}{k}$, rather than the usual $\binom{[n]}{k}$ (in other words, we identify elements of $[n]$ with points of V).

We proceed by contradiction. Suppose that there is a proper coloring of $\mathrm{KG}_{n,k}$ by at most $n-2k+1 = d+1$ colors. We fix one such proper coloring and we define sets $A_1, \ldots, A_{d+1} \subseteq S^d$: For a point $\boldsymbol{x} \in S^d$, we have $\boldsymbol{x} \in A_i$ iff there is at least one k-tuple $F \in \binom{V}{k}$ of color i contained in the open hemisphere centered at $\boldsymbol{x}$.

These sets $A_1, \ldots, A_{d+1}$ form an open cover of S^d, since each open hemisphere contains at least one k-tuple. By the Borsuk–Ulam theorem 2.1.1(3.2) (Lusternik–Schnirelmann for open covers), there exist $i \in [d+1]$ and $\boldsymbol{x} \in S^d$ such that $\boldsymbol{x}, -\boldsymbol{x} \in A_i$. In this way, we get two disjoint k-tuples colored by the color i, one in the open hemisphere centered at $\boldsymbol{x}$ and one in the opposite open hemisphere centered at $-\boldsymbol{x}$. This means that the considered coloring is not a proper coloring of the Kneser graph.

Proof of Gale's lemma. We prove the following version (equivalent to the above formulation using the central projection to S^d): there exist points $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_{2k+d}$ in $\mathbb{R}^{d+1}$ such that every open halfspace whose boundary hyperplane passes through $\mathbf{0}$ contains at least k of them.

The construction uses the moment curve (Definition 1.6.3) but we lift it one dimension higher, into the hyperplane $x_1 = 1$. That is, let

$$\bar{\gamma} := \{(1, t, t^2, \ldots, t^d) \in \mathbb{R}^{d+1} : t \in \mathbb{R}\}.$$

We take $2k+d$ arbitrary distinct points on $\bar{\gamma}$ and label them $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{2k+d}$ in the order in which they occur along the curve. For example, we can take $\boldsymbol{w}_i := \bar{\gamma}(i)$ for $1 \le i \le 2k+d$. We call the points $\boldsymbol{w}_2, \boldsymbol{w}_4, \ldots$ *even* and the points $\boldsymbol{w}_1, \boldsymbol{w}_3, \ldots$ *odd*. Further we define $\boldsymbol{v}_i := (-1)^i \boldsymbol{w}_i$.

Let h be a hyperplane passing through $\mathbf{0}$ and let h^+ and h^- be the two open halfspaces determined by it. We want to argue that both h^+ and h^- contain at least k points among the $\boldsymbol{v}_i$; we formulate the argument for h^+. Since $\boldsymbol{v}_i = \boldsymbol{w}_i$ for i even and $\boldsymbol{v}_i = -\boldsymbol{w}_i$ for i odd, we need to prove that *the number of even points $\boldsymbol{w}_i$ in h^+ plus the number of odd points $\boldsymbol{w}_i$ in h^- is at least k.*

Using Lemma 1.6.4, we see that every hyperplane h through the origin intersects $\bar{\gamma}$ at no more than d points. Moreover, if there are d intersections, then $\bar{\gamma}$ crosses h at each of the intersections.

Given an arbitrary h through the origin, we move it so that it contains the origin and exactly d points of $W := \{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{d+2k}\}$, while no point of W crosses from one side to the other during the motion. This is possible: having already $k < d$ points of W on h, we rotate h around the flat spanned by these points and $\mathbf{0}$, until we hit another point of W. After this motion, h intersects $\bar{\gamma}$ in exactly d points, which all lie in W.

Let W_{on} be the subset of the d points of W lying on h, and let $W_{\mathrm{off}} := W \setminus W_{\mathrm{on}}$ be the remaining $2k$ points. At every point of W_{on}, $\bar{\gamma}$ crosses from one side of h to the other.

Color a $\boldsymbol{w}_i \in W_{\mathrm{off}}$ black if either it is even and lies in h^+ or it is odd and lies in h^-. Otherwise, color $\boldsymbol{w}_i$ white. It is easy to see that as we follow $\bar{\gamma}$, black and while points of W_{off} alternate:

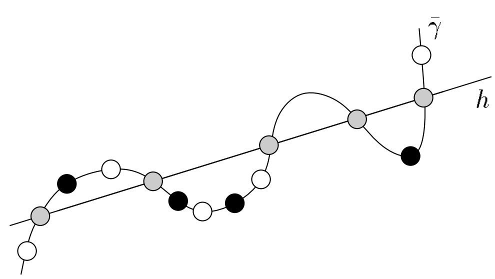

Indeed, let $\boldsymbol{w}$ and $\boldsymbol{w}'$ be two consecutive points of W_{off} along $\bar{\gamma}$ with j points of W_{on} between them. For j even, both $\boldsymbol{w}$ and $\boldsymbol{w}'$ are in the same halfspace and one of them is odd and the other is even, so one is black and one white. If j is odd, then $\boldsymbol{w}$ and $\boldsymbol{w}'$ are in different halfspaces but they are both even or both odd, and so again one is black and one white. So the number of black points is at least $\lfloor \frac{1}{2}|W_{\mathrm{off}}| \rfloor \ge k$. This proves Gale's lemma.

Schrijver's strengthening. Almost the same proof establishes a stronger theorem, first proved by Schrijver [Sch78] soon after Kneser's conjecture was proved.

Let us call a subset $S \in \binom{[n]}{k}$ *stable* if it does not contain any two adjacent elements modulo n; that is, if it corresponds to an independent set in the cycle C_n. We denote by $\binom{[n]}{k}_{stab}$ the family of stable k-subsets of $[n]$. The *Schrijver graph* $\mathrm{SG}_{n,k}$ is the induced subgraph of the Kneser graph $\mathrm{KG}_{n,k}$ on the stable k-sets. That is, the Schrijver graph $\mathrm{SG}_{n,k}$ has the stable k-subsets of $[n]$ as its vertices, and two vertices are connected by an edge if and only if they are disjoint k-sets.

Schrijver's theorem states that $\chi(\mathrm{SG}_{n,k}) = \chi(\mathrm{KG}_{n,k}) = n-2k+2$. In fact, Schrijver showed that $\mathrm{SG}_{n,k}$ is a *vertex-critical* subgraph of the Kneser graph $\mathrm{KG}_{n,k}$; that is, the chromatic number decreases by deleting any single vertex (stable k-set) from $\mathrm{SG}_{n,k}$; see Exercise 2.

The proof of Schrijver's theorem goes exactly as the one shown above for Kneser's conjecture, with the following strengthening of Gale's lemma: *there exists a $(2k+d)$-point set $V \subset S^d$ such that, under a suitable identification of V with $[n]$, every open hemisphere contains a stable k-tuple.* And this is exactly what the above proof of Gale's lemma provides: the black points form a stable set if the points of V are numbered along $\bar{\gamma}$.

Notes. Later we are going to present several more proofs of Kneser's conjecture. A summary of references and generalizations is given in the notes to Section 6.7.

Gale's proof of Lemma 3.3.3 is different from the one shown; it is more complicated and goes by induction on d and k. On the other hand, our argument is also based on Gale's work, namely on the investigation of cyclic polytopes, which are convex hulls of finite point sets on the moment curve. The possibility of proving both Gale's lemma and the stronger version needed for Schrijver's graphs by the above simple construction was observed by Ziegler.

Exercises

1. (a) Show that the fractional chromatic number of the Kneser graphs satisfies
$$\chi_f(\mathrm{KG}_{n,k}) \leq \frac{n}{k} \quad (n \geq 2k > 0).$$

   (b) Show that the inequality in (a) is actually an equality. Hint: (look up and) use the Erdős–Ko–Rado theorem.

2. (a) Show that the graph $\mathrm{SG}_{n,k}$ is vertex-critical (for chromatic number); that is, for every k-tuple $A \in V(\mathrm{SG}_{n,k})$, there is a proper coloring of the vertex set of $\mathrm{SG}_{n,k}$ by $n-2k+2$ colors that uses the color $n-2k+2$ only at A.

   (b) Show that not all $\mathrm{SG}_{n,k}$ are edge-critical (an edge may be removed without decreasing the chromatic number).

3. Show that the Schrijver graph $\mathrm{SG}_{n,k}$ is not regular in general; that is, its vertices need not all have the same degree. What can you say about the symmetries of the Schrijver graphs?

4. Show $\mathrm{KG}_{n,k}$ has no odd cycles of length shorter than $1+2\lceil\frac{k}{n-2k}\rceil$. What about even cycles?

3.4 Kneser's conjecture: second proof

The proof of Kneser's conjecture presented in this section is very natural and fairly simple. First we recall the important notion of the *chromatic number of a hypergraph* (or of a set system). If $\mathcal{S}$ is a system of subsets of a set X, a coloring $c: X \to [m]$ is a (proper) m-coloring of $(X, \mathcal{S})$ if no edge is monochromatic under c ($|c(S)| > 1$ for all $S \in \mathcal{S}$). The chromatic number $\chi(\mathcal{S})$ is the smallest m such that $(X, \mathcal{S})$ is m-colorable. In this section, we will only be interested in two-colorability.

Next, we define a less standard parameter of the set system $\mathcal{S}$: let the *m-colorability defect*, denoted by $\mathrm{cd}_m(\mathcal{S})$, be the minimum size of a subset $Y \subseteq X$ such that the system of the sets of $\mathcal{S}$ that contain no points of Y is m-colorable. In symbols,

$$\mathrm{cd}_m(\mathcal{S}) = \min\Big\{|Y| : (X \setminus Y, \{S \in \mathcal{S} : S \cap Y = \varnothing\}) \text{ is } m\text{-colorable}\Big\}.$$

For example, for $m = 2$, we want to color each point of X red, blue, or white in such a way that no set of $\mathcal{S}$ is completely red or completely blue (but it may be completely white), and $\mathrm{cd}_2(\mathcal{S})$ is the minimum required number of white points for such a coloring.

We prove

3.4.1 Theorem (Dol'nikov's theorem [Dol'81]). *For any finite set system $(X, \mathcal{S})$, we have*

$$\chi(\mathrm{KG}(\mathcal{S})) \geq \mathrm{cd}_2(\mathcal{S}).$$

Here the Kneser graph $\mathrm{KG}(\mathcal{S})$ is defined as in the previous section and χ is the usual graph-theoretic chromatic number.

If $\mathcal{S}$ consists of all the k-points subsets of $[n]$, $n \geq 2k$, then after deleting any at most $n-2k+1$ points, we are left with the system of all k-element subsets of a $(2k-1)$-element set. In any red-blue coloring of that set, one of the colors has at least k points and contains a monochromatic k-element set. Thus $\mathrm{cd}_2(\mathcal{S}) \geq n-2k+2$, and we see that Theorem 3.4.1 implies Kneser's conjecture.

For proving Theorem 3.4.1, we first need a geometric statement based on the Borsuk–Ulam theorem, slightly resembling the ham sandwich theorem.

3.4.2 Proposition. *Let $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_d$ be families of nonempty compact convex sets in $\mathbb{R}^d$, and suppose that for each $i = 1, 2, \ldots, d$, the system $\mathcal{C}_i$ is intersecting; that is, $C \cap C' \neq \varnothing$ for $C, C' \in \mathcal{C}_i$. Then there is a hyperplane (transversal) intersecting all the sets in $\bigcup_{i=1}^d \mathcal{C}_i$.*

Proof. For a direction vector $\boldsymbol{v} \in S^{d-1}$, let $\ell_{\boldsymbol{v}}$ denote the line containing $\boldsymbol{v}$ and passing through the origin, oriented from the origin towards $\boldsymbol{v}$. Consider the system of the orthogonal projections of the sets of $\mathcal{C}_i$ on the line $\ell_{\boldsymbol{v}}$:

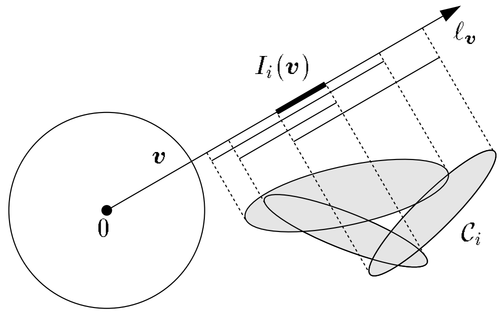

Each of these projections is a closed and bounded interval, and any two of them intersect. It is easy to see (directly, or by the one-dimensional Helly theorem) that the intersection of all these intervals is a nonempty interval, which we denote by $I_i(\boldsymbol{v})$. Let $\boldsymbol{m}_i(\boldsymbol{v})$ denote the midpoint of $I_i(\boldsymbol{v})$.

We define an antipodal mapping $f\colon S^{d-1} \to \mathbb{R}^d$, by letting $f(\boldsymbol{v})_i = \langle \boldsymbol{m}_i(\boldsymbol{v}), \boldsymbol{v} \rangle$ be the oriented distance of $\boldsymbol{m}_i(\boldsymbol{v})$ from the origin. This is a continuous antipodal map, and we claim that for any such map, there is a point $\boldsymbol{v} \in S^{d-1}$ with $f_1(\boldsymbol{v}) = f_2(\boldsymbol{v}) = \cdots = f_d(\boldsymbol{v})$. To see this, define a new antipodal map g, this time into $\mathbb{R}^{d-1}$, by letting $g_i = f_i - f_d$, $i = 1, 2, \ldots, d-1$. This g has a zero by the Borsuk–Ulam theorem, and if $g(\boldsymbol{v}) = \boldsymbol{0}$, then $f_1(\boldsymbol{v}) = f_2(\boldsymbol{v}) = \cdots = f_d(\boldsymbol{v})$ as required. For a $\boldsymbol{v}$ with this property, all the d midpoints $\boldsymbol{m}_i(\boldsymbol{v})$ coincide, and so the hyperplane passing through them and perpendicular to $\ell_{\boldsymbol{v}}$ is the desired transversal of all the sets in each $\mathcal{C}_i$.

Proof of Theorem 3.4.1. Suppose that there is a d-coloring of the Kneser graph $\mathrm{KG}(\mathcal{S})$. This means that $\mathcal{S}$ can be partitioned into set systems $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_d$, such that each two sets in $\mathcal{S}_i$ have a common point, $i = 1, 2, \ldots, d$.

Place the points of the ground set X into $\mathbb{R}^d$ in a general position, in such a way that no $d+1$ of them lie on a common hyperplane (and otherwise arbitrarily; for instance, they can be placed on the moment curve). Define the d families of convex sets in $\mathbb{R}^d$ by

$$\mathcal{C}_i = \{\mathrm{conv}(S) : S \in \mathcal{S}_i\}.$$

These $\mathcal{C}_i$ satisfy the assumptions of Proposition 3.4.2 above, and so there is a hyperplane h intersecting the convex hulls of all the $S \in \mathcal{S}$. Let $Y = X \cap h$ be the (at most d) points that lie on h:

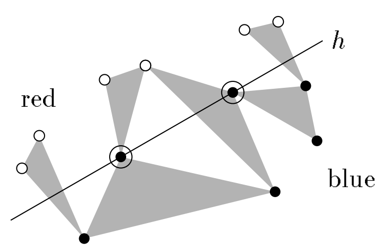

Color the points of $X \setminus Y$ in one of the open halfspaces of h red, and those in the other halfspace blue. Since none of the sets of S lies completely in one of the open halfspaces defined by h, this red-blue coloring shows that $\mathrm{cd}_2(\mathcal{S}) \leq d$. Theorem 3.4.1 is proved.

Notes. Theorem 3.4.1 is a special case of results of Dol'nikov [Dol'81] (also see [Dol'92], [Dol'94]). We postpone the discussion of his results and related material to the notes to Section 6.7. For another proof of Dol'nikov's theorem see Exercise 5.7.3.

Exercises

1. For set systems $\mathcal{S}$ with $\chi(\mathrm{KG}(\mathcal{S})) \leq 2$, prove Dol'nikov's theorem 3.4.1 by a direct combinatorial argument.

2. Find 2-colorable set systems $\mathcal{S}$ with $\chi(\mathrm{KG}(\mathcal{S}))$ arbitrarily large.

3. Show that for $n \geq 2k$, the family $\binom{[n]}{k}$ is *critical* with respect to the 2-colorability defect: If $\mathcal{S}$ is a proper subset of $\binom{[n]}{k}$, then
$$\mathrm{cd}_2(\mathcal{S}) \quad < \quad n - 2k + 2.$$

4. Show that the 2-colorability defect of the Schrijver hypergraphs $\binom{[n]}{k}_{stab}$ is given by
$$\mathrm{cd}_2(\tbinom{[n]}{k}_{stab}) \quad = \quad \min\{n, n - 4k + 4\}.$$

4

A Topological Interlude

In this chapter we explain some further basic topological concepts and constructions needed for the further development. We do it a little more thoroughly than necessary for our concrete applications. Similar to Chapter 1, most of the material should be well-known to readers fluent in elementary algebraic topology.

4.1 Quotient spaces

Given a topological space X and a subset $A \subset X$, we can form a new space by "shrinking A to a point." Two spaces can be "glued together" to form another space. A space can be factored using a group acting on it. All these important constructions are special cases of forming quotient spaces.

4.1.1 Definition (Quotient space). *Let X be a topological space and let $\approx$ be an equivalence relation on its elements. We define a topology on the set $X/\!\approx$ of equivalence classes as follows: A set $U \subseteq X/\!\approx$ is open if and only if $q^{-1}(U)$ is open in X, where $q\colon X \to X/\!\approx$ is the* **quotient map** *that maps each $x \in X$ to the equivalence class $[x]_\approx$ containing it.*

In constructions of quotient spaces, the equivalence $\approx$ is often given by a list of the nontrivial equivalence classes. That is, if $(A_i : i \in I)$ is some family of disjoint subsets of X, we define an equivalence $\approx$ on X corresponding to this family as follows: $x \approx y$ if and only if $x = y$ or there exists $i \in I$ with $x, y \in A_i$. Then we write $X/(A_i,\ i \in I)$ for $X/\!\approx$. The meaning is "the space $X/(A_i,\ i \in I)$ is obtained from X by shrinking each A_i to a single point." If we have only one $A_i = A$, we simply write X/A.

4.1.2 Example. Let $U = [0,1] \times [0,1]$ be the unit square. By gluing the two vertical sides together, i.e. by taking $U/(\{(0,y),(1,y)\}_{y\in[0,1]})$ we obtain the surface of a cylinder. The horizontal edges can be further glued either in a "direct" way (that is, a point $(x,0)$ is identified with $(x,1)$ for each $x \in [0,1]$), which produces a torus, or in a "twisted" way (i.e. a point $(x,0)$ is identified with $(1-x,1)$), which leads to the so-called Klein bottle (which cannot be embedded in $\mathbb{R}^3$, however).

Here are two other simple but useful constructions.

4.1.3 Definition (Sum and wedge). *Let X and Y be topological spaces.*

The **sum** *of X and Y, denoted by $X \sqcup Y$, corresponds to just "putting X and Y side by side." The point set of $X \sqcup Y$ is the disjoint union of X and Y (formally, we can take $(X \times \{0\}) \cup (Y \times \{1\})$, say) and each open set $U \subseteq X \sqcup Y$ is a (disjoint) union of an open set in X and an open set in Y.*

Now let $x_0 \in X$ and $y_0 \in Y$ some points (called **base points***). The* **wedge** *of X and Y, with respect to x_0 and y_0, is $X \vee Y := (X \sqcup Y)/(\{x_0, y_0\})$; that is, we take the sum and then glue x_0 to y_0.*

Many commonly encountered spaces (such as connected manifolds) are homogeneous, in the sense that for any $x, x' \in X$, there is a homeomorphism $h: X \to X$ with $h(x) = x'$. For such X, the choice of the base point in the wedge construction obviously doesn't matter.

The wedge is a special case of another construction: attaching one topological space to another by a given subspace, or gluing spaces.

4.1.4 Example (Gluing spaces). Let X and Y be topological subspaces with closed subspaces $A \subseteq X$ and $B \subseteq Y$ that are homeomorphic, with $h: A \to B$ being a given homeomorphism. The space obtained by gluing X and Y via h is the quotient space $X \sqcup_h Y$ obtained from the sum $X \sqcup Y$ as

$$X \sqcup_h Y := (X \sqcup Y)/\{\{a, h(a)\} : a \in A\}.$$

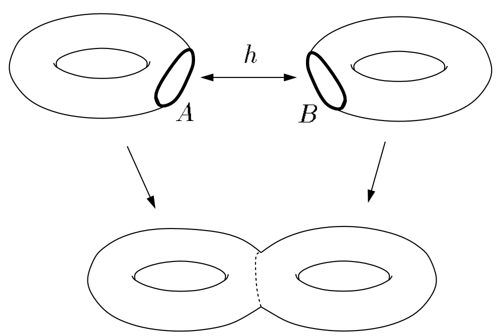

Our most significant instance of quotient spaces are joins, discussed in the next section. But first we mention a useful sufficient condition for homotopy equivalence.

4.1.5 Proposition (Contracting a contractible subcomplex is a homotopy equivalence). *Let X be the polyhedron of a simplicial complex K and $A \subseteq X$ the polyhedron of a subcomplex of K. Suppose that A is contractible. Then the quotient map $q: X \to X/A$ has a homotopy inverse; that is, a continuous map $p: X/A \to X$ such that $q \circ p \sim \mathrm{id}_{X/A}$ and $p \circ q \sim \mathrm{id}_X$. Therefore, $X \simeq X/A$.*

Many homotopy equivalences occurring "in practice" can be interpreted as sequences of operations according to Proposition 4.1.5 and their inverses. The conclusion holds for more general pairs (X, A) with A contractible; it is enough that they satisfy the "homotopy extension property" introduced in the proof below.

Proof. It is not entirely obvious how the required homotopy inverse p should be constructed; the reader may want to consider the example with $X = S^1$ and $A \subset X$ being a half-circle.

Let $(f_t: A \to A)_{t\in[0,1]}$ be a homotopy of the identity map $\mathrm{id}_A = f_0$ to the constant map f_1 with $f_1(a) = a_0 \in A$ for all $a \in A$. Suppose that we manage to extend this homotopy to a homotopy $(\bar{f}_t)_{t\in[0,1]}$ on the whole X, with $\bar{f}_0 = \mathrm{id}_X$ (each $\bar{f}_t: X \to X$ coincides with f_t on A). Then $\bar{f}_1$ is a continuous map $X \to X$ that is constant on A, and so we can consider it as a map $p: X/A \to X$ (formally, $p([x]) = \bar{f}_1(x)$ for $x \in X$). We have $p(q(x)) = p([x]) = \bar{f}_1(x)$, and so $(\bar{f}_t)_{t\in[0,1]}$ is a homotopy witnessing $p \circ q \sim \mathrm{id}_X$. As for the other direction, we note that if we set $p_t([x]) = [\bar{f}_t(x)]$, we obtain well-defined maps (since each $\bar{f}_t$ maps A into A), which provide a homotopy of $p_0 = \mathrm{id}_{X/A}$ with $p_1 = q \circ p$ as required.

It remains to show that the homotopy can indeed be extended. This is a special case of the following definition:

4.1.6 Definition. *Let X be a topological space and $A \subseteq X$ a subspace of it. We say that the pair (X, A) has the* **homotopy extension property** *if every continuous mapping $F: (A\times[0,1])\cup(X\times\{0\}) \to Y$, where Y is some topological space, can be extended to a continuous mapping $\bar{F}: X \times [0,1] \to Y$:*

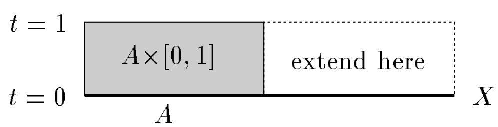

In our case, we have the homotopy $(f_t: A \to A)_{t\in[0,1]}$ and an extension $\bar{f}_0: X \to X$ of f_0. So we set

$$F(x,t) := \begin{cases} f_t(x) & \text{for } x \in A, t \neq 0 \\ \bar{f}_0(x) & \text{for } x \in X, t = 0. \end{cases}$$

For the proof of Proposition 4.1.5, it thus suffices to show that whenever X is the polyhedron of a (finite) simplicial complex K and A is the polyhedron of a subcomplex of K, then the pair (X, A) has the homotopy extension property. (Here we do not need contractibility of A anymore; this was used for the existence of the homotopy $\mathrm{id}_A \sim \text{const.}$)

To establish the homotopy extension property of a pair (X, A), it is enough to verify that $S := (A\times[0,1])\cup(X\times\{0\})$ is a deformation retract of $T := X\times[0,1]$. Indeed, if $(g_t)_{t\in[0,1]}$ is a deformation retraction witnessing this, we simply set $\bar{F}(z) = F(g_1(z))$, $z = (x,t) \in X\times[0,1]$. This works since $g_1(z) \in S$ for all $z \in T$ and $g_1(z) = z$ on S.[1]

[1] We haven't used the full power of deformation retraction, only the existence of a single continuous map g_1 with the two properties just stated. The existence of such g_1 defines the weaker concept of S being a *retract* of T.

The deformation retraction of T on S is constructed gradually. First we note that the deformation retraction exists if X is a simplex and A is its boundary, as the picture indicates for a 1-dimensional simplex:

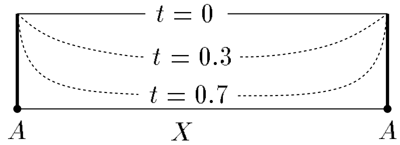

(In Exercise 2, the reader is invited to construct such a "hollowing out" deformation retraction explicitly.) Then we hollow out the simplices of X not lying in A one by one, starting with those of the largest dimension and proceeding to the smaller dimensions, until only the simplices of A remain "fat."

Exercises

1. Check that $B^d/(S^{d-1}) \cong S^d$.

2. Let σ be a (geometric) simplex. Describe a deformation retraction of $\sigma\times[0,1]$ to $(\partial\sigma \times [0,1]) \cup (\sigma\times\{0\})$, either geometrically or by an explicit formula.

3. Let K be a simplicial complex and $\mathsf{K}_1, \mathsf{K}_2 \subseteq \mathsf{K}$ subcomplexes that together cover K (i.e. $\mathsf{K} = \mathsf{K}_1 \cup \mathsf{K}_2$). Assume that both K_2 and $\mathsf{K}_1 \cap \mathsf{K}_2$ are contractible. Using Proposition 4.1.5, prove that $\mathsf{K} \simeq \mathsf{K}_1$; in particular, if K_1 is contractible, then K is contractible as well.

4. Consider a finite graph G as a 1-dimensional simplicial complex (the vertices of the graph are the vertices of the simplicial complex and the edges are the 1-dimensional simplices). Suppose that G is connected and has n vertices and m edges. Show that G is homotopy equivalent to a wedge of $m-n+1$ circles (S^1's).

5. Let X be a space and let $A \subseteq X$ be such that the pair (X, A) has the homotopy extension property. Let Y be another space, let $B, C \subseteq Y$, and let $h\colon A \to B$ and $g\colon A \to C$ be homeomorphisms such that h and g are homotopic as maps $X \to Y$. Prove that $X \sqcup_h Y$ and $X \sqcup_g Y$ are homotopy equivalent.

4.2 Joins (and products)

For many mathematical structures, including topological spaces, we have a notion of a Cartesian product $X \times Y$. For topological spaces, $X \times Y$ has the set-theoretical Cartesian product of X and Y as the set of points, and the topology of $X \times Y$ is the coarsest one making the projections maps $\pi_X\colon X \times Y \to X$ and $\pi_Y\colon X \times Y \to Y$ continuous. More explicitly, the topology on $X \times Y$ is generated by the "open rectangles" $U \times V$, where $U \subseteq X$ and $V \subseteq Y$ are open sets.

When working with simplicial complexes, a drawback of the Cartesian product is that the product of two simplices is not a simplex, except for trivial cases:

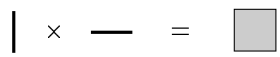

So if we want to regard a product of simplicial complexes as a simplicial complex, we have to triangulate it. We now introduce another product-like operation on topological spaces called *join* and denoted by $*$. The first advantage over the Cartesian product is that the join of simplices is again a simplex:

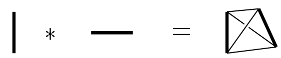

Other advantages are subtler and we will encounter some of them later.

The join has an extremely natural combinatorial definition for simplicial complexes.

4.2.1 Definition (Join of simplicial complexes). *Let K and L be simplicial complexes. Assuming that $V(\mathsf{K}) \cap V(\mathsf{L}) = \emptyset$, the join $\mathsf{K} * \mathsf{L}$ has vertex set $V(\mathsf{K}) \cup V(\mathsf{L})$ and simplices $F \cup G$, for all $F \in \mathsf{K}$ and all $G \in \mathsf{L}$.*

If the vertex sets are not disjoint, we formally rename the vertices so that they become disjoint. So we let $V(\mathsf{K} * \mathsf{L}) := (V(\mathsf{K}) \times \{0\}) \cup (V(\mathsf{L}) \times \{1\})$ and the simplices are $F * G := (F \times \{0\}) \cup (G \times \{1\})$ for all $F \in \mathsf{K}$ and all $G \in \mathsf{L}$.

For simplices we have $\sigma^k * \sigma^\ell = \sigma^{k+\ell+1}$ and $(\sigma^0)^{*n} = \sigma^{n-1}$; here σ^0 is a single point and K^{*n} means the n-fold join $\mathsf{K} * \mathsf{K} * \cdots * \mathsf{K}$. Note that K^{*n} has $n|V(\mathsf{K})|$ vertices, with one copy of $V(\mathsf{K})$ for each factor.

4.2.2 Example (important!). As a more interesting and challenging example, we consider $(S^0)^{*n}$, where S^0 is the 0-dimensional sphere consisting of two isolated vertices; call them a and b. The n-fold join has $2n$ vertices $a_1, a_2, \ldots, a_n$ and $b_1, b_2, \ldots, b_n$. A subset of this vertex set is a simplex in $(S^0)^{*n}$ if and only if it does not contain both a_i and b_i for some i. Comparing with the description of the proper faces of the crosspolytope below Definition 1.4.1, we get that $(S^0)^{*n}$ is the surface of the n-dimensional crosspolytope, i.e. a triangulation of S^{n-1}. We conclude

$$\|(S^0)^{*n}\| \cong S^{n-1}.$$

Since the join is obviously associative, we further get $S^k * S^\ell \cong S^{k+\ell+1}$ (considering the crosspolytope triangulations on both sides).

The join is also defined for arbitrary topological spaces:

4.2.3 Definition (Join of spaces). *Let X and Y be topological spaces. The join $X * Y$ is the quotient space $X \times Y \times [0,1]/\approx$, where the equivalence relation $\approx$ is given by $(x, y, 0) \approx (x', y, 0)$ for all $x, x' \in X$ and all $y \in Y$ ("for $t = 0$, x does not matter") and $(x, y, 1) \approx (x, y', 1)$ for all $x \in X$ and all $y, y' \in Y$ ("for $t = 1$, y does not matter").*

The drawing illustrates this definition for X and Y being line segments (1-simplices):

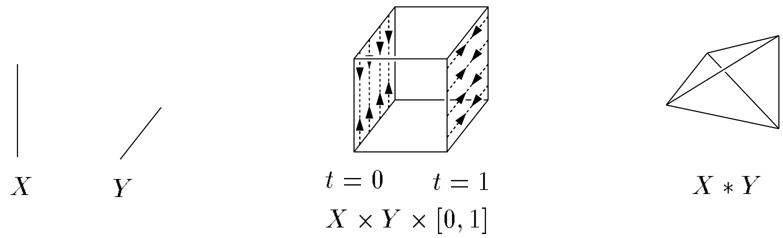

Here is a helpful geometric interpretation of the join:

4.2.4 Proposition (Geometric join). *Suppose that X and Y are subspaces of some Euclidean space, and that $X \subseteq U$ and $Y \subseteq V$, where U and V skew affine suspaces of some $\mathbb{R}^n$ (that is, $U \cap V = \emptyset$ and the affine hull of $U \cup V$ has dimension $\dim U + \dim V + 1$). Moreover, suppose that both X and Y are bounded. Then the space*

$$Z := \{t\boldsymbol{x} + (1-t)\boldsymbol{y} : t \in [0,1], \boldsymbol{x} \in X, \boldsymbol{y} \in Y\} \subset \mathbb{R}^n,$$

i.e. the union of all segments connecting a point of X to a point of Y, is homeomorphic to the join $X * Y$.

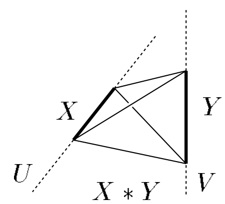

Sketch of proof. There is an obvious continous map

$$X \times Y \times [0,1] \to \{t\boldsymbol{x} + (1-t)\boldsymbol{y} : t \in [0,1], \boldsymbol{x} \in X, \boldsymbol{y} \in Y\}$$

that induces a homeomorphism

$$(X \times Y \times [0,1])/\approx \; \to \{t\boldsymbol{x} + (1-t)\boldsymbol{y} : t \in [0,1], \boldsymbol{x} \in X, \boldsymbol{y} \in Y\}.$$

First we observe that $t'\boldsymbol{x}' + (1-t')\boldsymbol{y}' = t''\boldsymbol{x}'' + (1-t'')\boldsymbol{y}''$ implies $t' = t''$, and, if $t \neq 0$, also $\boldsymbol{x} = \boldsymbol{x}'$. From it follows that our map is a bijection. The continuity at points with $t \neq 0$ is fairly obvious. For $t \in \{0,1\}$, some care is needed, and one needs to use the boundedness of X and Y (for unbounded X and Y, the inverse mapping need *not* be continuous).

With this interpretation, it is not hard to see the equivalence of the definition of join for simplicial complexes with that for spaces; that is, $\|\mathsf{K} * \mathsf{L}\| \cong \|\mathsf{K}\| * \|\mathsf{L}\|$ for any simplicial complexes K and L. Indeed, it suffices to check that if X is a k-simplex and Y is an ℓ-simplex, the geometric definition in Proposition 4.2.4

yields a $(k+\ell+1)$-simplex. This follows since if $A \subset U$ and $B \subset V$ are affinely independent sets, then $A \cup B$ is affinely independent, too.

Yet another description of the join is presented in Exercise 3.

The join is commutative, in the sense $X * Y \cong Y * X$. It is also associative, $(X * Y) * Z \cong X * (Y * Z)$, as is best seen from the definition for simplicial complexes (at least for triangulable spaces).

Cone and suspension. Two other well-known topological constructions can be seen as special cases of the join. The *cone* over a space X is the join with a one-point space: $\mathrm{cone}(X) := X * \{p\}$. Geometrically, the cone is the union of all segments connecting the points of X to a new point. Another equivalent definition is the quotient space $\mathrm{cone}(X) \cong (X \times [0,1])/(X \times \{1\})$:

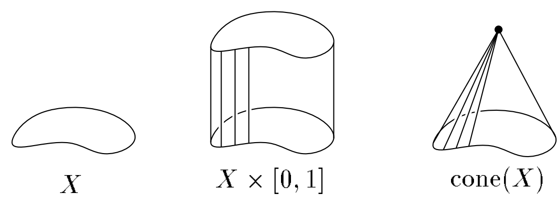

The join with a two-point space, $X * S^0$, is called the *suspension* of X and denoted by $\mathrm{susp}(X)$. It can be interpreted as erecting a double cone over X, or as the quotient $(X \times [0,1])/(X \times \{0\}, X \times \{1\})$.

Notation for points of a join. Let us consider an n-fold join X^{*n}. A point in it can be conveniently written in the form ("formal convex combination") $t_1 x_1 + t_2 x_2 + \cdots + t_n x_n$, where $t_1, t_2, \ldots, t_n$ are nonnegative reals summing up to 1 and $x_1, x_2, \ldots, x_n$ are points of X. As the notation suggests, if $t_i = 0$, then the choice of x_i does not matter, and we get the same point of X^{*n} for all $x_i \in X$. On the other hand, the analogy with convex combination should not be pushed too far: this formal convex combination is *not commutative*; for example, $\frac{1}{2}a + \frac{1}{2}b$, $a \neq b$, is a point of X^{*2} different from $\frac{1}{2}b + \frac{1}{2}a$. This is because of the "renaming convention" for joins: we should really think of x_1 as coming from a different copy of X than x_2, and so on.

Join as a functor. Joins can be naturally defined not only for spaces but also for (continuous) maps. Given maps $f\colon X_1 \to X_2$ and $g\colon Y_1 \to Y_2$, a map $f * g\colon X_1 * Y_1 \to X_2 * Y_2$ is given by $tx + (1-t)y \mapsto tf(x) + (1-t)g(y)$.

Joins and products. The Cartesian product $X \times Y$ can be embedded into $X * Y$ by $(x, y) \mapsto \frac{1}{2}x + \frac{1}{2}y \in X * Y$. Similarly, the Cartesian power X^n embeds into X^{*n} by $(x_1, x_2, \ldots, x_n) \mapsto \frac{1}{n}x_1 + \frac{1}{n}x_2 + \cdots + \frac{1}{n}x_n$. Here is an illustration for our usual example $X = Y = \sigma^1$:

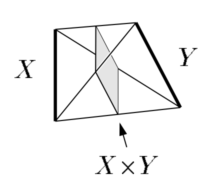

Exercises

1. Verify the following homeomorphisms and homotopy equivalences (X and Y are triangulable spaces). If you cannot do the general case in (d)–(f), try at least some special cases like $X = Y = S^1$.

   (a) $\mathrm{cone}(S^n) \cong B^{n+1}$,

   (b) $\mathrm{cone}(B^n) \cong B^{n+1}$,

   (c) $\mathrm{susp}(B^n) \cong B^{n+1}$,

   (d) $\mathrm{susp}(X \vee Y) \simeq \mathrm{susp}(X) \vee \mathrm{susp}(Y)$,

   (e) $\mathrm{susp}(X \sqcup Y) \simeq \mathrm{susp}(X) \vee \mathrm{susp}(Y) \vee S^1$,

   (f) $\mathrm{susp}((X \vee Y) \sqcup \{p\}) \simeq \mathrm{susp}(X) \vee \mathrm{susp}(Y) \vee S^1$.

   Parts (d)–(f) may fail if X and Y are arbitrary topological spaces.

2. Show that joins preserve homotopy equivalence; that is, if $X \simeq X'$, then $X * Y \simeq X' * Y$.

3. (Another interpretation of the join) Let X and Y be spaces. Verify that $X * Y$ is homeomorphic to the subspace $(\mathrm{cone}(X) \times Y) \cup (X \times \mathrm{cone}(Y))$ of the product $\mathrm{cone}(X) \times \mathrm{cone}(Y)$. (Equivalently, glue the two spaces $\mathrm{cone}(X) \times Y$ and $X \times \mathrm{cone}(Y)$ in the subspaces homeomorphic to $X \times Y$ that are given by the inclusions of bases $X \subseteq \mathrm{cone}(X)$ and $Y \subseteq \mathrm{cone}(Y)$.)

4. Let the topology on a space X be induced by a metric ρ and the topology on Y by a metric σ. Assume that both ρ and σ are bounded, i.e. no two points have distance more than K for a suitable fixed number K. Construct a metric τ on the join $X * Y$ inducing its topology (and check that it indeed works). Warning: there are some quite tempting wrong solutions.

5. In Section 1.7, we associated the simplicial complex $\Delta(P)$ with every (finite) poset P. What is the appropriate operation "$*$" on posets, such that $\Delta(P * Q) = \Delta(P) * \Delta(Q)$?

4.3 *k*-connectedness

Informally, a topological space X is k-connected if it has no "holes" up to dimension k. A hole in dimension ℓ is something that prevents some suitably placed S^ℓ from continuously shrinking to a point:

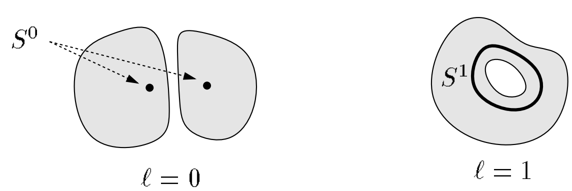

(To make a hole in a B^3 in dimension 0, slice it in two pieces; for dimension 1, puncture a tunnel in it, and for dimension 2, make a void inside.) Of course, things can be more complicated: a torus certainly has a hole in dimension 1 in this sense, but what about dimension 2? Fortunately, we need not contemplate such finesses here, since the formal definition is simple:

4.3.1 Definition (k-connected space). *Let $k \geq -1$. A topological space X is k-connected if for every $\ell = -1, 0, 1, \ldots, k$, each continuous map $f\colon S^\ell \to X$ can be extended to a continuous map $\bar{f}\colon B^{\ell+1} \to X$. (Equivalently, each $f\colon S^\ell \to X$ is nullhomotopic.)*

Here S^{-1} is interpreted as $\emptyset$ and B^0 as a single point, and so (-1)-connected means nonempty.

For $k \geq 0$, k-connectedness includes the condition (for $\ell = 0$) that X has to be arcwise connected. A space X satisfying the condition for $\ell = 1$, i.e. with every map $S^1 \to X$ nullhomotopic, is usually called *simply connected.* So 1-connected means arcwise connected and simply connected.

It is not hard to check that homotopy equivalence preserves k-connectedness (Exercise 1). Another very believable result is

4.3.2 Theorem. *The n-sphere S^n is $(n-1)$-connected and not n-connected.*

Proof. By the Borsuk-Ulam Theorem 2.1.1 (1.3), S^n is not n-connected.

The fact that S^n is $(n-1)$-connected may seem almost obvious, but one has to be careful, as already maps $S^1 \to S^n$ can be quite wild (think of a space-filling curve!).

Let us consider a (uniformly) continuous map $f\colon S^k \to S^n$. We show that it is homotopic to another map $g\colon S^k \to S^n$ that is not surjective. Such a g is obviously nullhomotopic and hence f, too, is nullhomotopic.

To construct g, we find an $\varepsilon > 0$ such that $\|f(\boldsymbol{x}) - f(\boldsymbol{y})\| < 1$ whenever $\|\boldsymbol{x} - \boldsymbol{y}\| < \varepsilon$, and a triangulation Δ of S^k such that every simplex in Δ has diameter smaller than ε (think of S^k as the boundary of a simplex, for example). Now we define g on each simplex $\sigma \in \Delta$ by interpolating the values of f at the vertices of σ suitably. Moreover, such a definition yields a homotopy of f and g. Namely, we define $F\colon S^k \times I \to S^n$ by

$$F(\boldsymbol{x}, t) := \frac{t\sum_{i=1}^{m} \lambda_i f(\boldsymbol{v}_i) + (1-t) f(\boldsymbol{x})}{\|t\sum_{i=1}^{m} \lambda_i f(\boldsymbol{v}_i) + (1-t) f(\boldsymbol{x})\|}$$

where $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m$ are the vertices of $\operatorname{supp}(\boldsymbol{x})$ (the simplex of Δ containing $\boldsymbol{x}$ in its relative interior) and $\boldsymbol{x} = \sum_{i=1}^{m} \lambda_i \boldsymbol{v}_i$. We need to show that the denominator is never 0. All the $f(\boldsymbol{v}_i)$, as well as $f(\boldsymbol{x})$, have distance at most 1 from $\boldsymbol{v}_1$ and hence they all lie in a spherical cap of radius smaller than 1. So their convex hull cannot contain the origin and F is well-defined and continuous. We have $f = F(*, 0)$ and we set $g := F(*, 1)$. For every simplex $\sigma \in \Delta$, the image $g(\sigma)$ is contained in a hyperplane in $\mathbb{R}^{n+1}$ passing through the origin. A finite union of hyperplanes cannot cover the sphere and hence g is not surjective.

In many topological proofs of geometric or combinatorial results, the problem is reduced to showing that certain spaces are highly connected. Number of tools are available for the latter task. We will soon explain a simple trick (Sarkaria's inequality) which will allow us to avoid explicit proofs of k-connectedness in most of the applications. But for attacking other problems, it can be useful to have tools for establishing k-connectedness at hand. In the rest of this section, we state some such results without proof (since they use a technical apparatus which we do not want to assume in this book). Later we will add a few more, which we will be able to prove even with our meager topological means.

Homology and k-connectedness. The following theorem refers to the reduced singular homology groups. A reader not familiar with homology may just want to know that they are parameters of a topological space, invariant under homotopy equivalence and efficiently computable for simplicial complexes (and for many other spaces).

4.3.3 Theorem. *Let X be a nonempty topological space and let $k \geq 1$. Then X is k-connected if and only if it is simply connected (i.e. the fundamental group $\pi(X) = 0$) and $\tilde{H}_i(X) = 0$ for all $i = 0, 1, \ldots, k$.*

This is a special case of a famous theorem of Hurewicz: for a simply connected space, the first nonzero homotopy and homology groups occur in the same dimension and they are isomorphic; see e.g. Hatcher [Hat01].

Since the kth homology group of a simplicial complex depends only on simplices of dimension at most $k{+}1$, and the fundamental group depends only on the 2-skeleton, we have

4.3.4 Proposition. *A simplicial complex K is k-connected if and only if the $(k{+}1)$-skeleton $\mathsf{K}^{\leq k+1}$ is k-connected.*

This can also be proved directly, without resorting to homology. Another consequence of Theorem 4.3.3 (and of formulas for the homology of a join), which does not seem easy to prove directly, is

4.3.5 Proposition (Connectivity of join). *Suppose that X is k-connected and Y is ℓ-connected, where X and Y are triangulable (or CW-complexes). Then $X * Y$ is $(k{+}\ell{+}2)$-connected.*

It may also be useful to know that if a k-dimensional simplicial complex is k-connected, then it is contractible. This follows, for example, from a theorem of Whitehead and Theorem 4.3.3. (For general spaces this need not be true!) Moreover, finite k-dimensional $(k{-}1)$-connected simplicial complexes have a special structure: they are homotopy equivalent to a point or to a wedge of k-dimensional spheres.

Exercises

1. Prove that if X is k-connected and $Y \simeq X$, then Y is k-connected as well.
2. (a) Suppose that X is a space that is not k-connected. Show that $X \times Y$ cannot be k-connected either, for any Y.

   (b) Prove that if both X and Y are k-connected, then so is $X \times Y$.
3. (a) Deduce from Theorem 4.3.2 that $S^n \not\simeq S^m$ unless $m = n$.

   (b) Use (a) to derive $\mathbb{R}^n \not\cong \mathbb{R}^m$ unless $m = n$.

4.4 Cell complexes

This section is optional: cell complexes are generally nice and very useful in topology, they will be mentioned in the formulation of some of the subsequent general theorems, but they will not be essential for any of our concrete applications.

In algebraic topology, cell complexes are usually called *CW-complexes*. The meaning of the mysterious letters C and W will be explained soon, but right now we note that they are significant only for complexes with infinitely many cells. We will occasionally use the name *cell complex* for a finite CW-complex.

Informally, a CW-complex is a topological space that can be pasted together from finite-dimensional balls, where a new k-ball is always glued by its boundary to the part already made from balls of dimension $< k$. Thus, we start with a discrete set of vertices, called the 0-*cells* in this context. Then we put in some 1-balls, called 1-*cells*. A 1-cell is just a closed interval, whose two endpoints are glued to some vertices, possibly both to the same vertex. The spaces obtained at this phase can be viewed as topological realizations of graphs, possibly with loops and multiple edges:

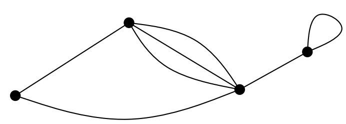

Next, we can paste in some 2-dimensional discs (2-cells). The boundary of each disc is glued to some of the edges, possibly in a complicated manner. Here are a few examples of what can be obtained with a single 2-cell. We can make the disc (as a topological space) with one 0-cell, one 1-cell, and one 2-cell:

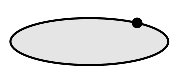

With just one 0-cell, *no* 1-cell, and one 2-cell, we can manufacture an S^2; note that the whole boundary of the 2-cell is identified to a point:

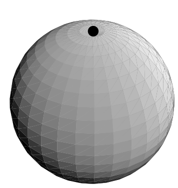

Of course, an S^2 can be made in many other ways, too; for example, using 2 cells of each dimension 0,1,2, as will be shown in a drawing in Section 5.2. If we picture a 2-cell as a square and we paste the edges in the indicated manner to two 1-cells a and b, we get a torus:

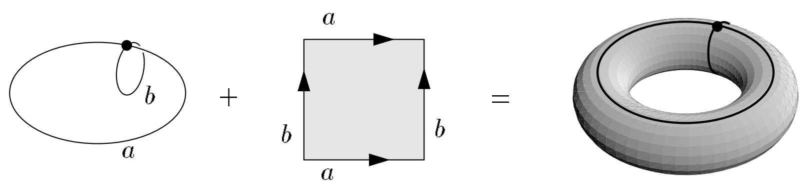

In fact, as is taught in basic courses of algebraic topology (such as [Mun00] or [Sti93]), we can get *any* 2-dimensional manifold without boundary, including non-orientable ones like the projective plane or the Klein bottle, from a regular convex polygon by suitable boundary identifications.

A (geometric) simplicial complex is a special case of a CW-complex (each simplex is homeomorphic to a ball). One obvious new thing in CW-complexes is that, while simplices are "straight," cells can be "curved." But another, perhaps less obvious difference is that a simplex must remain homeomorphic to a ball in the simplicial complex, including the boundary, while the boundary of a cell may become glued to itself and entangled in a complicated manner. For example, it is legal to glue a 2-cell to the middle of a 1-cell:

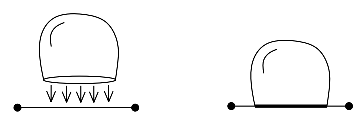

Here is a formal definition of a CW-complex. A *CW-complex* is a Hausdorff space X which is the union of a collection $\{e_\alpha\}_{\alpha\in\Lambda}$ of disjoint subspaces called *cells* with the following properties.

- Each e_α has some *dimension* $\dim e_\alpha \in \{0,1,2,\ldots\}$. The *n-skeleton* of X is
$$X^{\le n} = \bigcup\{e_\alpha : \alpha\in\Lambda, \dim e_\alpha \le n\}.$$
- If $\dim e_\alpha = n$, then there is a continuous *characteristic map* (or *attachment map*) $\chi_\alpha\colon B^n \to X$, such that $\partial B^n = S^{n-1}$ is mapped into the $(n-1)$-skeleton $X^{\le n-1}$ and $\operatorname{int} B^n$ is mapped homeomorphically onto e_α.

These conditions are sufficient to define a finite CW-complex (i.e. one with finitely many cells); the topology on X is determined uniquely by the characteristic maps. Note that a finite CW-complex is always compact. An infinite CW-complex has to satisfy the following two additional conditions (which are automatically satisfied by finite CW-complexes):

- (Weak topology) A set $F \subseteq X$ is closed if and only if $F \cap \bar{e}_\alpha$ is closed for each $\alpha\in\Lambda$, where $\bar{e}_\alpha$ denotes $\chi_\alpha(B^n)$, i.e. the cell e_α together with its boundary.

- (Closure finiteness) The boundary of each cell e_α, i.e. the image of ∂B^n under χ_α, intersects only finitely many cells.

The "morphisms" of CW-complexes are called cellular maps. A map $f: X \to Y$ of CW-complexes is *cellular* if, for each $n \geq 0$, the n-skeleton $X^{\leq n}$ is mapped into the n-skeleton $Y^{\leq n}$. If a cellular map is a homeomorphism, then any n-cell is mapped homeomorphically onto an n-cell.

For many applications, a CW-complex structure for a space is as good as a triangulation, or nearly as good. At the same time, the CW-complex structure can have just a couple of cells where a triangulation would have to be quite large. For instance, an S^n can be expressed as a cell complex with one 0-cell and one n-cell (as we have seen for S^2), while the smallest triangulation is the boundary of an $(n+1)$-simplex, with $2^{n+1} - 1$ simplices!

Although there exist non-triangulable CW-complexes, it is known that every CW-complex X is homotopy equivalent to a polyhedron of a simplicial complex K. Moreover, one may assume $\dim \mathsf{K} = \dim X$ and if X is finite, then K can be chosen finite as well.

A *subcomplex* of a CW-complex X is a closed subspace $A \subseteq X$ that is the union of some of the cells of X (recall that the cells are relatively open). A nice feature of CW-complexes, not shared by simplicial complexes, is that the quotient X/A is again a CW-complex (Exercise 1).

If A is a subcomplex of a CW-complex X, then the pair (X, A) has the homotopy extension property; this is proved almost exactly as for simplicial complexes. Proposition 4.1.5 also extends without any difficulty: if A is contractible, then $X/A \simeq X$.

Notes. There are several restricted classes of CW-complexes that lie between general CW-complexes and simplicial complexes.

In a *regular* (finite) cell complex, we require that each of the attachment maps χ_α be a homeomorphism (not only on the interior of B^n but also on the boundary). The intersection of the boundaries of two closed cells can still be topologically nontrivial, but regular cell complexes admit a simple combinatorial description. Namely, if we define the partial order on the set of closed cells by inclusion, then the order complex of this poset is homeomorphic to the original cell complex (and it is natural to call the resulting simplicial complex the *first barycentric subdivision* of the regular cell complex).

A more special class of regular cell complexes are *polyhedral* complexes. At least two different definitions appear in the literature. A more strict definition is very similar to the definition of a simplicial complex but the cells can be convex polytopes, instead of just simplices. Every two cells intersect in a cell and a face of a cell is again a cell. In a more permissive definition, it is only required that every cell be homeomorphic to a convex polytope.

A special case of the latter definition are the Δ-*complexes*, used in [Hat01]. The cells are simplices, they are still glued together face-to-face, but for example gluing two triangles by just two sides is permitted:

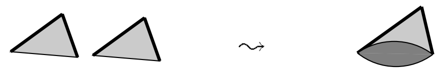

A (geometric) Δ-complex is obtained from a family of disjoint simplices by face identifications. More precisely, let $(\sigma_\alpha : \alpha \in A)$ be a family of (geometric) simplices. We assume that for each σ_α, some linear ordering of the vertices has been fixed. Further let $(\mathcal{F}_\beta \beta \in B)$ be given, where each $\mathcal{F}_\beta$ is a family of simplices, all simplices in $\mathcal{F}_\beta$ having the same dimension k_β and each of them being a face of some σ_α. The Δ-complex specified by these data is obtained from the sum $\bigsqcup_{\alpha \in A} \sigma_\alpha$ by identifying all the faces in each $\mathcal{F}_\beta$ to a single k_β-face. The identification is made according to the canonical affine homeomorphisms among the faces in $\mathcal{F}_\beta$ that extend the (unique) order-preserving bijections of the vertex sets. Note that F_β may contain several faces of the same σ_α; so, for example, the three edges of a triangle can all be identified as indicated by the arrows:

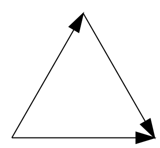

(The resulting mind-boggling geometric object can be realized in $\mathbb{R}^3$ and it is known as the *dunce cap*.) Unlike general CW-complexes, the specification of a Δ-complex is purely combinatorial, albeit formally more complicated than for a simplicial complex. Let us remark that modern homotopy theory uses yet another generalization of simplicial complexes, called the *simplicial sets*; these are always infinite and at present they do not seem relevant for combinatorial applications in the spirit discussed here.

Exercises

1. Let X be a CW-complex and A a subcomplex of it. Define a cell structure on X/A and check that it is a CW-complex (if you like, assume that X is finite).

5

Nonexistence of $\mathbb{Z}_2$-Maps

In the applications covered in Chapter 3, we always associated a continuous map of the sphere with the considered problem, sometimes in a quite natural way (for the ham sandwich cut theorem, say) and sometimes by a clever ad hoc construction (in both the proofs of Kneser's conjecture). The Borsuk–Ulam theorem applied to this map then provided the desired object or a contradiction.

Here we first generalize the Borsuk–Ulam theorem from spheres to a much wider class of spaces, which gives us more flexibility. We pursue just one among many possible directions of generalizations, dealing with the $\mathbb{Z}_2$-index, which proved very fruitful in combinatorial and geometric applications. Then we introduce constructions, most notably deleted joins, which for many problems lead to a suitable space with a continuous map in an almost canonical way. In this connection, one speaks about a *configuration space* (encoding all possible "configurations" in the considered problem) and a *test map* (distinguishing configurations with some desired property from the others, say by mapping them to zero).

5.1 $\mathbb{Z}_2$-spaces and $\mathbb{Z}_2$-maps

One of the versions of the Borsuk–Ulam theorem asserts that there is no antipodal map $S^{n+1} \to S^n$, and this is the starting point of our generalizations. We will view antipodal maps not only as maps between topological spaces, but rather as maps between topological spaces with additional structure given by the antipodality. Thus, here we regard S^n as the pair $(S^n, -)$, where "$-$" is a shorthand for the mapping $\boldsymbol{x} \mapsto -\boldsymbol{x}$. The antipodality "$-$" is a homeomorphism of the underlying space (S^n, or also $\mathbb{R}^n$), and it gives the identity if performed twice: $-(-\boldsymbol{x}) = \boldsymbol{x}$. These are the essential properties that are reflected in the definition of a general "antipodality space." Anticipating the terminology of the subsequent generalizations, we begin to use brave new names for old things, though: we start saying $\mathbb{Z}_2$-*action* instead of antipodality and $\mathbb{Z}_2$-*map* instead of antipodal map.

5.1.1 Definition ($\mathbb{Z}_2$-space and $\mathbb{Z}_2$-map). *A* **$\mathbb{Z}_2$-space** *is a pair (X, ν), where X is a topological space and $\nu\colon X \to X$ is a homeomorphism, called the* **$\mathbb{Z}_2$-action** *on X, such that $\nu^2 = \nu \circ \nu = \mathrm{id}_X$.*

The $\mathbb{Z}_2$-action ν is **free** *if $\nu(x) \neq x$ for all $x \in X$; that is, if ν has no fixed point. In that case, the $\mathbb{Z}_2$-space (X, ν) is also called free.*

If (X, ν) and (Y, ω) are $\mathbb{Z}_2$-spaces, a **$\mathbb{Z}_2$-map** *$f\colon (X, \nu) \to (Y, \omega)$ is a continuous map $X \to Y$ that commutes with the $\mathbb{Z}_2$-actions: for all $x \in X$, we have $f(\nu(x)) = \omega(f(x))$, or, more briefly, $f \circ \nu = \omega \circ f$.*

A $\mathbb{Z}_2$-map is also called an *equivariant map*, or an *involution*, or an *antipodal map*. If the $\mathbb{Z}_2$-action on a $\mathbb{Z}_2$-space (X, ν) is understood, we write just "$\mathbb{Z}_2$-space X;" this is similar to the conventions for many other mathematical structures.

Obvious examples of $\mathbb{Z}_2$-spaces are $(S^n, -)$ and $(\mathbb{R}^n, -)$. Here is one example that, in reality, is not very different, but at least it looks different at first sight.

5.1.2 Example. Consider the boundary of the $(n-1)$-dimensional simplex as an abstract simplicial complex K; i.e. $\mathsf{K} = 2^{[n]} \setminus \{[n]\}$. Let $\mathsf{L} = \mathrm{sd}(\mathsf{K})$ be the first barycentric subdivision of K; thus, the vertex set of L consists of all proper nonempty subsets of $[n]$. Define a simplicial map $\nu\colon V(\mathsf{L}) \to V(\mathsf{L})$ by setting, for a vertex $F \in V(\mathsf{L})$, $\nu(F) = [n] \setminus F$. A simplex in L is a chain of sets under inclusion, and so ν maps simplices to simplices (reversing the inclusion!). Moreover, ν is surjective (all chains are obtained) and $\nu^2 = \mathrm{id}$. So $(\|\mathsf{L}\|, \|\nu\|)$ is a (free) $\mathbb{Z}_2$-space.

As we know, $\|\mathsf{L}\| = \|\mathsf{K}\| \cong S^{n-2}$. For $n = 3$, the action ν is depicted below:

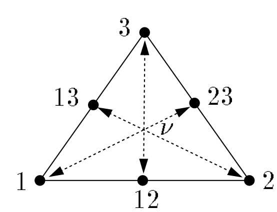

It is essentially the same as the usual antipodality "$-$" on S^1. (As we will see later, all free $\mathbb{Z}_2$-actions on S^n are equivalent for our purposes.)

The L above is an example of a simplicial $\mathbb{Z}_2$-complex. In general, a *simplicial $\mathbb{Z}_2$-complex* is a simplicial complex K with a simplicial map $\nu\colon V(\mathsf{K}) \to V(\mathsf{K})$ such that $\|\nu\|$ is a $\mathbb{Z}_2$-action on $\|K\|$. A cell $\mathbb{Z}_2$-complex is defined analogously: it is a finite CW-complex and the $\mathbb{Z}_2$-action is a cellular map.

5.1.3 Example (Join of $\mathbb{Z}_2$-spaces). If (X_1, ν_1) and (X_2, ν_2) are $\mathbb{Z}_2$-spaces, the join $X_1 * X_2$ can be equipped with the $\mathbb{Z}_2$-action $\nu_1 * \nu_2$. The join of free $\mathbb{Z}_2$-spaces is clearly free.

The two-point space S^0 has an obvious free $\mathbb{Z}_2$-action that exchanges the two points (in the standard embedding of S^0 into $\mathbb{R}^1$, it is precisely the usual action $\boldsymbol{x} \mapsto -\boldsymbol{x}$). As we saw in Example 4.2.2, the n-fold join $(S^0)^{*n}$ is an S^{n-1}, namely the boundary of the n-dimensional crosspolytope. By considering this join as a $\mathbb{Z}_2$-space, we recover the standard $\mathbb{Z}_2$-action $\boldsymbol{x} \mapsto -\boldsymbol{x}$ on the boundary of the crosspolytope.

The next examples look rather simple, but we will be encountering their variations all the time.

5.1.4 Example ($\mathbb{Z}_2$-action on $X \times X$). Let X be any space. The Cartesian product $X \times X$ can be made into a $\mathbb{Z}_2$-space by letting the $\mathbb{Z}_2$-action exchange the two components; $\nu\colon (x, y) \mapsto (y, x)$.

5.1.5 Example ($\mathbb{Z}_2$-action on $X * X$). Similarly, the join $X * X$ becomes a $\mathbb{Z}_2$-space if we define the $\mathbb{Z}_2$-action ν by $tx + (1-t)y \mapsto (1-t)y + tx$ (recall the convention about writing the points in a join as formal convex combinations, and visualize this action for X being a segment).

The $\mathbb{Z}_2$-spaces in the last two examples are not free. Later on, we will be using constructions making them free by deleting suitable points from $X \times X$ or from $X * X$.

Exercises

1. Verify that the $\mathbb{Z}_2$-action ν in Example 5.1.2 is indeed free.

5.2 The $\mathbb{Z}_2$-index

Let (X, ν) and (Y, ω) be $\mathbb{Z}_2$-spaces. Let us write $X \xrightarrow{\mathbb{Z}_2} Y$ if there exists a $\mathbb{Z}_2$-map from X to Y and $X \overset{\mathbb{Z}_2}{\not\longrightarrow} Y$ if no $\mathbb{Z}_2$-map exists. The Borsuk–Ulam theorem tells us that $S^{n+1} \overset{\mathbb{Z}_2}{\not\longrightarrow} S^n$. In the applications of the concepts developed in this chapter, the crux is always in showing $X \overset{\mathbb{Z}_2}{\not\longrightarrow} Y$ for some given X and Y. Of course, the relation $\xrightarrow{\mathbb{Z}_2}$ is rather complicated and one should not expect to be able to decide whether $X \xrightarrow{\mathbb{Z}_2} Y$ for arbitrary given X and Y (in view of the difficulty of both homeomorphism and homotopy equivalence, for example). Nevertheless, with the tools introduced later one can succeed in many interesting concrete cases.

The relation $\xrightarrow{\mathbb{Z}_2}$ is obviously transitive, and it is useful to think of it as a partial ordering: if $X \xrightarrow{\mathbb{Z}_2} Y$, then Y is at least as big as X. To support this ideology notationally, we also write

$$X \le_{\mathbb{Z}_2} Y \quad \text{if} \quad X \xrightarrow{\mathbb{Z}_2} Y.$$

Strictly speaking, $\le_{\mathbb{Z}_2}$ is not a partial ordering but rather a partial quasiordering, since many spaces are equivalent under it (homeomorphic spaces with "the same" $\mathbb{Z}_2$-actions, for example).

Before proceeding, we can observe that non-free $\mathbb{Z}_2$-spaces are uninteresting from the point of view of $\le_{\mathbb{Z}_2}$. Namely, if (Y, ω) is such that $\omega(y_0) = y_0$, then $X \xrightarrow{\mathbb{Z}_2} Y$ for all X: simply send all of X to y_0. In the $\le_{\mathbb{Z}_2}$ relation, all non-free $\mathbb{Z}_2$-spaces are equivalent and strictly larger than all free $\mathbb{Z}_2$-spaces.

The $\mathbb{Z}_2$-index. Spheres were useful in the Borsuk–Ulam theorem, and here we are going to use them as a yardstick for measuring the "size" of $\mathbb{Z}_2$-spaces with respect to $\le_{\mathbb{Z}_2}$.

5.2.1 Definition ($\mathbb{Z}_2$-index). *Let (X, ν) be a $\mathbb{Z}_2$-space. We set*

$$\mathrm{ind}_{\mathbb{Z}_2}(X) := \min\{n \in \{0,1,2,\ldots\} : X \xrightarrow{\mathbb{Z}_2} S^n\}.$$

Here S^n is taken with the standard antipodal action.

The $\mathbb{Z}_2$-index can be a natural number or ∞; the latter happens, for example, for a non-free $\mathbb{Z}_2$-space.

5.2.2 Proposition (Properties of the $\mathbb{Z}_2$-index).

(i) *If $X \leq_{\mathbb{Z}_2} Y$, then $\mathrm{ind}_{\mathbb{Z}_2}(X) \leq \mathrm{ind}_{\mathbb{Z}_2}(Y)$. Therefore, $\mathrm{ind}_{\mathbb{Z}_2}(X) > \mathrm{ind}_{\mathbb{Z}_2}(Y)$ implies $X \overset{\mathbb{Z}_2}{\not\longrightarrow} Y$.*

(ii) *$\mathrm{ind}_{\mathbb{Z}_2}(S^n) = n$, for all $n \geq 0$ (with the standard $\mathbb{Z}_2$-action on S^n).*

(iii) *$\mathrm{ind}_{\mathbb{Z}_2}(X * Y) \leq \mathrm{ind}_{\mathbb{Z}_2}(X) + \mathrm{ind}_{\mathbb{Z}_2}(Y) + 1$.*

(iv) *If X is $(n-1)$-connected, then $\mathrm{ind}_{\mathbb{Z}_2}(X) \geq n$.*

(v) *If X is a* **free** *simplicial $\mathbb{Z}_2$-complex (or cell $\mathbb{Z}_2$-complex) of dimension n, then $\mathrm{ind}_{\mathbb{Z}_2}(X) \leq n$.*[1]

Part (i) follows trivially from the definition (right?) and it suggests how the $\mathbb{Z}_2$-index can be used for establishing the nonexistence of a $\mathbb{Z}_2$-map. The condition $\mathrm{ind}_{\mathbb{Z}_2}(X) > \mathrm{ind}_{\mathbb{Z}_2}(Y)$ is only *sufficient* for $X \overset{\mathbb{Z}_2}{\not\longrightarrow} Y$. If $\mathrm{ind}_{\mathbb{Z}_2}(X) \leq \mathrm{ind}_{\mathbb{Z}_2}(Y)$, both the possibilities $X \xrightarrow{\mathbb{Z}_2} Y$ and $X \overset{\mathbb{Z}_2}{\not\longrightarrow} Y$ are still open, although examples of the second possibility are not obvious (see the notes and Exercise 4).

Part (ii) is essentially a version of the Borsuk–Ulam theorem.

Part (iii) follows immediately from $S^n * S^m \cong S^{n+m+1}$. As we will see, it can sometimes be used to show that the $\mathbb{Z}_2$-index of some space is *large*, in the form $\mathrm{ind}_{\mathbb{Z}_2}(X) \geq \mathrm{ind}_{\mathbb{Z}_2}(X * Y) - \mathrm{ind}_{\mathbb{Z}_2}(Y) - 1$ for a suitable Y.

Finally, parts (iv) and (v) are a little more difficult and we prove them below. The statement (iv), $\mathrm{ind}_{\mathbb{Z}_2}(X) \geq n$ for $(n-1)$-connected X, is the basic tool for bounding the $\mathbb{Z}_2$-index below, while (v), $\mathrm{ind}_{\mathbb{Z}_2}(X) \leq \dim(X)$, is typically used to bound it above.

Mapping the sphere: proof of (iv). To show that $\mathrm{ind}_{\mathbb{Z}_2}(X) \geq n$ for an $(n-1)$-connected X, it suffices to exhibit a $\mathbb{Z}_2$-map $g: S^n \to X$. We proceed

[1] With some more technical machinery, this claim can be shown for much more general spaces X. Namely, if X is paracompact, then $\mathrm{ind}_{\mathbb{Z}_2}(X) \leq \dim(X)$. Paracompactness is a mild topological condition satisfied by practically all the usually encountered topological spaces; for example, by all metric spaces. A topological space X is paracompact if it is Hausdorff and each open cover $\mathcal{U}$ of X has a locally finite open refinement $\mathcal{V}$. Here a cover $\mathcal{U}$ is *open* if it consists of open sets, a cover $\mathcal{V}$ is a *refinement* of a cover $\mathcal{U}$ if each set of $\mathcal{V}$ is contained in some set of $\mathcal{U}$, and $\mathcal{V}$ is *locally finite* if each point of X has an open neighborhood intersecting only finitely many members of $\mathcal{V}$.

The dimension is the usual *covering dimension*. For a metric space X, $\dim X \leq n$ if every finite open cover of X has a finite open refinement such that each point of X is contained in at most $n+1$ sets of the refinement. For a detailed treatment of both paracompactness and topological dimensions see [Eng77]. For finite simplicial complexes and CW-complexes, the covering dimension coincides with the maximum dimension of a simplex or cell, respectively.

by induction, constructing $\mathbb{Z}_2$-maps $g_k: S^k \to X$ by induction on k. The cases $k = -1$ and $k = 0$ are clear. For the induction step, consider S^{k-1} as a subset of S^k, by identifying it with the "equator" $\{\boldsymbol{x} \in S^k: x_{k+1} = 0\}$. Furthermore, via the projection map $\pi: \mathbb{R}^{k+1} \to \mathbb{R}^k$ that deletes the last coordinate, the upper hemisphere $S^k_+ := \{\boldsymbol{x} \in S^k: x_{k+1} \geq 0\}$ is homeomorphic to the ball B^k. Now if a $\mathbb{Z}_2$-map $g_{k-1}: S^{k-1} \to X$ has been constructed, we can extend it to a continuous map $\bar{g}_{k-1}: B^k \to X$, since X is $(k-1)$-connected. Using π we can then define g_k on S^k_+, as

$$g_k = \bar{g}_{k-1} \circ \pi: \ S^k_+ \ \longrightarrow \ B^k \ \longrightarrow \ X.$$

Setting $g_k(\boldsymbol{x}) := \nu(g_k(-\boldsymbol{x}))$ for $\boldsymbol{x} \in S^k_-$ (the lower hemisphere), we get a map $g_k: S^k \to X$. This map is well-defined since g_k is antipodal on the intersection S^{k-1} of the two hemispheres of S^k. It is continuous since it is continuous on both the closed hemispheres of S^k, and it is a $\mathbb{Z}_2$-map by construction.

It is instructive to unwrap this inductive proof; for concreteness, we do it for $n = 2$. First we regard S^0 as two antipodal points S^0_+ and S^0_- in $\mathbb{R}^3$. We choose the value at S^0_+ as an arbitrary $x_0 \in X$, and the value at S^0_- is enforced: $\nu(x_0)$. Next, we extend to an arc S^1_+ connecting S^0_+ and S^0_-, using the 0-connectedness of X, and we again put the enforced values on the opposite arc S^1_-. The two arcs combine to a full circle S^1, and from this circle, we extend to the upper hemisphere S^2_+ by the 1-connectedness of X. We finish the construction by assigning the antipodal values on the lower hemisphere.

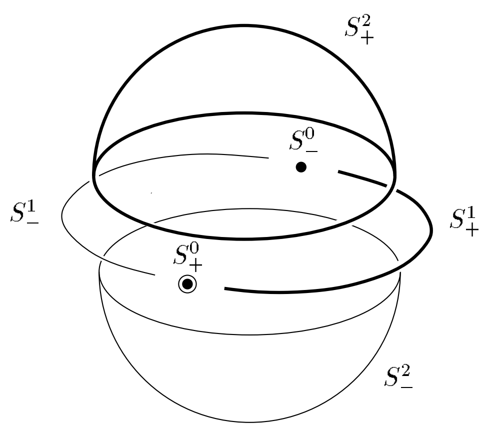

The proof implicitly used a suitable cell decomposition of S^2 (see Section 4.4). This decomposition is *equivariant*, meaning that the interior of each k-dimensional cell is mapped bijectively onto the interior of another k-dimensional cell by the $\mathbb{Z}_2$-action.

In order to stay in the realm of the perhaps more familiar simplicial complexes, we can also do the proof using an antipodally symmetric triangulation of S^k. For example, in the usual octahedral triangulation of S^2, we can choose the values of g_0 at the three marked vertices,

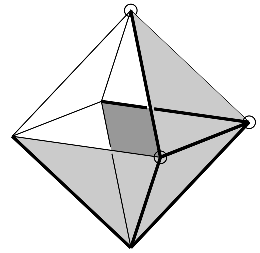

get the values at the other vertices by antipodality, extend on the marked edges, and so on.

Mapping into the sphere: proof of (v). Here it suffices to construct a $\mathbb{Z}_2$-map $g\colon \|\mathsf{K}\| \to S^n$ for every free simplicial $\mathbb{Z}_2$-complex with $\dim \mathsf{K} \leq n$. We show that, more generally, a free n-dimensional simplicial $\mathbb{Z}_2$-complex can be $\mathbb{Z}_2$-mapped into any $(n-1)$-connected $\mathbb{Z}_2$-space Y. The argument is almost exactly as in the previous proof.

We construct $\mathbb{Z}_2$-maps $g_k\colon \|\mathsf{K}^{\leq k}\| \to Y$ by induction, $k = 0, 1, \ldots, n$. Having already constructed g_k, we divide the $(k+1)$-dimensional simplices in K into equivalence classes—the orbits under the $\mathbb{Z}_2$-action; one can check that each class consists of two disjoint simplices F and $\nu(F)$ (Exercise 1). We pick one simplex from each class and for these simplices, we extend g_k on the interior using the k-connectedness of Y. We then define g_{k+1} on the interiors of the remaining simplices in the only possible way that makes g_{k+1} a $\mathbb{Z}_2$-map. The same proof goes through for cell $\mathbb{Z}_2$-complexes.

Other $\mathbb{Z}_2$-indices. There are various other sensible ways of defining a "$\mathbb{Z}_2$-index;" the one we have used is technically quite simple but others may be more powerful or easier to compute in some cases. In principle, any mapping from the class of $\mathbb{Z}_2$-spaces to some partially ordered set that is monotone with respect to the ordering $\leq_{\mathbb{Z}_2}$ can serve as a "$\mathbb{Z}_2$-index." But in order to get interesting results, the mapping should satisfy some extra properties similar to (ii)–(v) in Proposition 5.2.2. A little about other notions of index will be mentioned below and in the notes to Section 6.2.

Notes. A parameter of a $\mathbb{Z}_2$-space X called the *genus* and equal, in our notation, to $1+\mathrm{ind}_{\mathbb{Z}_2}(X)$, was introduced by Krasnosel'skiĭ [Kra52].

Another, similar (but not always equivalent) notion of index for $\mathbb{Z}_2$-spaces was defined by Yang [Yan54]; his definition can be expressed using a suitable equivariant homology theory with $\mathbb{Z}_2$-coefficients. He proved that if (X, ν) is a $\mathbb{Z}_2$-space of index n (in particular, if $(X, \nu) = (S^n, -))$ and $f\colon X \to \mathbb{R}^m$ is a continuous map, then (his) index of the coincidence set $A_f = \{x \in X : f(x) = f(\nu(x)\}$ is at least $n-m$, and consequently, $\dim(A_f) \geq n-m$, too. He derived generalizations of several nice geometric theorems listed below. Some of these results were obtained by Bourgin [Bou55], too.

Kakutani-type results. Kakutani [Kak43] proved that for any compact convex set in $\mathbb{R}^3$ there exists a cube circumscribed to it and touching it by all the 6 facets. This is an easy consequence of the following: for any continuous $f: S^2 \to \mathbb{R}$, there are three mutually perpendicular vectors $\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3 \in S^2$ with $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = f(\boldsymbol{x}_3)$. This was generalized to dimension n (with $n+1$ mutually orthogonal vectors) by Yamabe and Yujobô [YY50], and re-derived by Yang [Yan54] (in a greater generality, with an arbitrary $\mathbb{Z}_2$-space of index n replacing S^n, with a suitable abstract notion of "orthogonality"). Yang [Yan54] and Bourgin [Bou63] proved that for any continuous $f: S^n \to \mathbb{R}$, there are n mutually orthogonal $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \in S^n$ with $f(\boldsymbol{x}_1) = f(-\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = \cdots = f(-\boldsymbol{x}_n)$, generalizing such result for S^2 due to Dyson [Dys51]. Here is another nice result of Yang of this type: if $f: S^{mn+m+n} \to \mathbb{R}^m$ is continuous, then there exists an antipodally symmetric subset of S^{mn+m+n} of index (and dimension) at least n on which f is constant. Numerous results about circumscribed geometric shapes and similar problems can be found in the work of Makeev, such as [Mak96].

In this connection, we should also mention a conjecture of Knaster [Kna47], stating that for any continuous $f: S^n \to \mathbb{R}^m$ and any configuration $K \subset S^n$ of $n-m+2$ points, there exists a rotation ρ of S^n such that $f(\rho(K))$ is a single point. Although this was proved for some special values of n and m and for some special configurations, the general conjecture was eventually refuted by Makeev [Mak84]. Stronger counterexamples were provided by Babenko and Bogatyĭ [BB89]; for example, they showed that if $n-2 \leq mt$, then there is a mapping $f: S^n \to \mathbb{R}^m$, given by a polynomial of degree at most t, such that no configuration of $2t+1$ points on a great circle in S^n has the property required by Knaster's conjecture.

A CW-complex with $S^3 \stackrel{\mathbb{Z}_2}{\not\to} X \stackrel{\mathbb{Z}_2}{\not\to} S^2$. The example we are going to sketch was constructed with the help of R. Živaljević and P. Csorba; we do not give a full proof. Let $h: S^3 \to S^2$ be the Hopf map (see [Hat01] or other topology textbooks). Construct X by attaching two 4-cells (copies of B^4) to the standard S^2, where the boundary of the first cell is attached by h and the boundary of the other cell by $-h$. The $\mathbb{Z}_2$-action ν acts on the S^2 as the antipodality and it interchanges the two 4-cells. If there were a $\mathbb{Z}_2$-map $S^3 \to X$, it could be deformed so that it remains a $\mathbb{Z}_2$-map and goes into the 3-skeleton of X. But the 3-skeleton is just the S^2, and so such map doesn't exist. If $f: S^2 \to S^2$ is a $\mathbb{Z}_2$-map, it can be shown that $f \circ h: S^3 \to S^2$ is not nullhomotopic (using the properties of the Hopf invariant), and so it cannot be extended to a map $B^4 \to S^2$. But a $\mathbb{Z}_2$-map $X \to S^2$ would yield such an extension.

Exercises

1. Let K be a simplicial complex and let ν be a free simplicial $\mathbb{Z}_2$-action on K. Prove that $F \cap \nu(F) = \varnothing$ for every $F \in \mathsf{K}$.

2. Give examples of free $\mathbb{Z}_2$-spaces of index n that are not $(n-1)$-connected.

3. Give an example of a free $\mathbb{Z}_2$-space X with $\mathrm{ind}_{\mathbb{Z}_2}(X) = \infty$.

4. Define the following index-like quantity for a $\mathbb{Z}_2$-space X:

$$\operatorname{dni}(X) := \max\{n \geq 0 : S^n \xrightarrow{\mathbb{Z}_2} X\}.$$

(a) Formulate and prove analogues of Proposition 5.2.2(i)–(v) for $\operatorname{dni}(X)$, and check that $\operatorname{dni}(X) \leq \operatorname{ind}(X)$ for all $\mathbb{Z}_2$-spaces X.

(b) Call a free $\mathbb{Z}_2$-space X *tidy* if $\operatorname{dni}(X) = \operatorname{ind}_{\mathbb{Z}_2}(X) < \infty$. Show that if X and Y are tidy, then $X \xrightarrow{\mathbb{Z}_2} Y$ if and only if $\operatorname{ind}_{\mathbb{Z}_2}(X) \leq \operatorname{ind}_{\mathbb{Z}_2}(Y)$.

(c) Construct an example of a free $\mathbb{Z}_2$-space X with $\operatorname{dni}(X) = 0 < \operatorname{ind}_{\mathbb{Z}_2}(X)$ (in particular, X is not tidy).

5.3 The topological Radon theorem

Many proofs concerning geometric embeddability, coloring of Kneser-like graphs, and other applications of topological methods have a common general scheme. In this section we encounter it for the first time.

We begin with a result well-known in convex geometry.

5.3.1 Theorem (Radon's theorem). *Every set $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{d+2}\}$ of $d+2$ points in $\mathbb{R}^d$ can be divided into two disjoint subsets whose convex hulls intersect.*

It may be good practice to visualize this for $d \leq 2$. For $d = 1$ we have three points on the real line, $x_1 \leq x_2 \leq x_3$, say. Then $\{x_2\}$ intersects $[x_1, x_3]$. For $d = 2$, four points are given in the plane. Then either one point $\boldsymbol{x}_i$ is contained in the convex hull of the others, and then we have the partition into $\{\boldsymbol{x}_i\}$ and $X \setminus \{\boldsymbol{x}_i\}$, or the four points form the vertices of a convex quadrilateral, and then the diagonals are the two intersecting convex hulls.

Although the standard proof is simple and unrelated to topology, we outline it for completeness.

Proof. Any $d+2$ points in $\mathbb{R}^d$ are affinely dependent. Let us fix an affine dependence: $\alpha_1 \boldsymbol{x}_1 + \alpha_2 \boldsymbol{x}_2 + \cdots + \alpha_{d+2} \boldsymbol{x}_{d+2} = \mathbf{0}$, $\sum_{i=1}^{d+2} \alpha_i = 0$. Then we define $I_1 := \{i \in [d+2] : \alpha_i > 0\}$ and $I_2 := [d+2] \setminus I_1$. Further let $S := \sum_{i \in I_1} \alpha_i = \sum_{j \in I_2} (-\alpha_j)$. Then the point $\boldsymbol{x} := \sum_{i \in I_1} (\frac{\alpha_i}{S}) \boldsymbol{x}_i = \sum_{j \in I_2} (-\frac{\alpha_j}{S}) \boldsymbol{x}_j$ is a convex combination of points in $X_1 := \{\boldsymbol{x}_i : i \in I_1\}$ as well as a convex combination of points in $X_2 = X \setminus X_1$.

An equivalent formulation of Radon's theorem. *For every affine map $f\colon \|\sigma^{d+1}\| \to \mathbb{R}^d$ there exist two disjoint faces F_1, F_2 of the $(d+1)$-simplex σ^{d+1} such that $f(\|F_1\|) \cap f(\|F_2\|) \neq \varnothing$.*

Proof of the equivalence. Each such f is determined by the images of the $d+2$ vertices of the simplex. The image of a face is the convex hull of the images of its vertices(Exercise 3).

We prove a significant generalization of Radon's theorem, which shows that very little of the vector-space structure of $\mathbb{R}^d$ is needed for the validity of Radon's theorem.

5.3.2 Theorem (Topological Radon's theorem; Bajmóczy & Bárány [BB79]). *Let $f\colon \|\sigma^{d+1}\| \to \mathbb{R}^d$ be a continuous map. Then there exist two disjoint faces F_1, F_2 of σ^{d+1} such that $f(\|F_1\|) \cap f(\|F_2\|) \neq \emptyset$.*

Since this is a prototype of several similar but more complicated statements to come later, it is important to realize what it asserts. For example, for $d = 1$, it says that if the triangle is mapped into the line, there are some two disjoint faces, typically a side of the triangle and its opposite vertex, whose images intersect. For $d = 2$, we have a tetrahedron mapped into the plane, and the theorem tells us that again the images of some two disjoint faces intersect; in this case, they can be a triangle and its opposite vertex or two opposite edges.

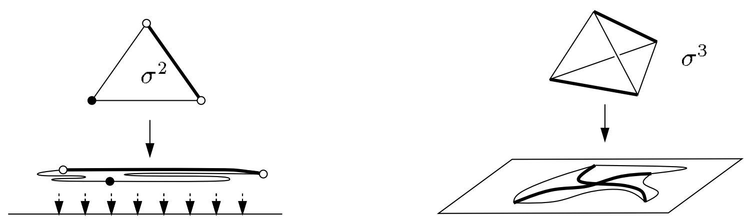

If we recall the notion of support of a point $\boldsymbol{x}$ in a geometric simplicial complex (the simplex containing $\boldsymbol{x}$ in its relative interior), we can also express the theorem by saying that there are $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \|\sigma^{d+1}\|$ with disjoint supports and such that $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2)$.

The first key idea in the proof is to pass to Cartesian products. Namely, let

$$\begin{aligned} f_2\colon\ \|\sigma^{d+1}\| \times \|\sigma^{d+1}\| &\to \mathbb{R}^d \times \mathbb{R}^d \\ (\boldsymbol{x}_1, \boldsymbol{x}_2) &\mapsto (f(\boldsymbol{x}_1), f(\boldsymbol{x}_2)). \end{aligned}$$

The theorem can now be reformulated as follows: there is a pair $\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2) \in \|\sigma^{d+1}\|^2$ such that $\operatorname{supp}(\boldsymbol{x}_1) \cap \operatorname{supp}(\boldsymbol{x}_2) = \emptyset$ and $f_2(\boldsymbol{x})$ intersects the *diagonal* in $(\mathbb{R}^d)^2$, i.e. the set $\{(\boldsymbol{y}, \boldsymbol{y}) : \boldsymbol{y} \in \mathbb{R}^d\}$.

For contradiction, let us suppose that this is not the case. This means that f_2 maps the set

$$X := \{(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \|\sigma^{d+1}\|^2 : \operatorname{supp}(\boldsymbol{x}_1) \cap \operatorname{supp}(\boldsymbol{x}_2) = \emptyset\}$$

into the set

$$Y := \{(\boldsymbol{y}_1, \boldsymbol{y}_2) \in (\mathbb{R}^d)^2 : \boldsymbol{y}_1 \neq \boldsymbol{y}_2\}.$$

The next crucial observation is that both X and Y are *free $\mathbb{Z}_2$-spaces*. Indeed, if we define $\nu\colon X \to X$ by $\nu((\boldsymbol{x}_1, \boldsymbol{x}_2)) = (\boldsymbol{x}_2, \boldsymbol{x}_1)$, then ν is obviously a homeomorphism with $\nu^2 = \mathrm{id}_X$ and, moreover, $\nu((\boldsymbol{x}_1, \boldsymbol{x}_2)) \neq (\boldsymbol{x}_1, \boldsymbol{x}_2)$ since X contains no pairs $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ with $\boldsymbol{x}_1 = \boldsymbol{x}_2$. Similarly, we define $\omega\colon Y \to Y$ by $\omega((\boldsymbol{y}_1, \boldsymbol{y}_2)) = (\boldsymbol{y}_2, \boldsymbol{y}_1)$, and we check in exactly the same way that ω is a free $\mathbb{Z}_2$-action on Y. Finally, f_2 is a $\mathbb{Z}_2$-map (it was carefully constructed that way). So, after these tricks, namely introducing the Cartesian products and deleting suitable points from them, we are in the situation for which we have been preparing in the previous sections: we have specific $\mathbb{Z}_2$-spaces X and Y and we would like to prove $X \stackrel{\mathbb{Z}_2}{\not\to} Y$. If we manage to do so, we reach a contradiction with the original assumption, namely that the images of disjoint faces of σ^{d+1} under f never intersect.

Deleted products. Before we start inspecting our specific X and Y, let us mention a general terminology. If Z is a space, the *(twofold) deleted product* of Z, denoted by Z^2_Δ, is the space

$$Z^2_\Delta := (Z \times Z) \setminus \{(x, x) : x \in Z\}.$$

So in our situation above, we have $Y = (\mathbb{R}^d)^2_\Delta$.

Our X is also a kind of deleted product, but this time we delete more: the product of each simplex with itself. If Δ is a geometric simplicial complex, we define its deleted product:

$$\Delta^2_\Delta := \{\sigma_1 \times \sigma_2 : \sigma_1, \sigma_2 \in \Delta,\ \sigma_1 \cap \sigma_2 = \varnothing\}.$$

It can be checked that this is a polyhedral cell complex. Moreover, its polyhedron (i.e. the union of its cells) is determined by the underlying abstract simplicial complex of Δ up to homeomorphism. So for an abstract simplicial complex K, the topological space corresponding to its deleted product is well-defined, and we denote it by $\|\mathsf{K}^2_\Delta\|$. (Note that $\|\mathsf{K}^2_\Delta\|$ is typically *not* homeomorphic to $\|\mathsf{K}\|^2_\Delta$, although it can be shows that they are homotopy equivalent and have the same $\mathbb{Z}_2$-index.) We can also write $\|\mathsf{K}^2_\Delta\| = \{(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \|\mathsf{K}\|^2 \colon \mathrm{supp}(\boldsymbol{x}_1) \cap \mathrm{supp}(\boldsymbol{x}_2) = \varnothing\}$. In our case, we have $X = \|(\sigma^{d+1})^2_\Delta\|$.

Both Z^2_Δ and $\|\mathsf{K}^2_\Delta\|$ can be made into free $\mathbb{Z}_2$-spaces as above, the actions being the exchanges of coordinates.

What are our deleted products? Now that we have set the stage, for concluding the proof of the topological Radon's theorem it would suffice to show $X \stackrel{\mathbb{Z}_2}{\not\to} Y$. To this end, as we know, it would be enough to prove $\mathrm{ind}_{\mathbb{Z}_2}(X) > \mathrm{ind}_{\mathbb{Z}_2}(Y)$.

It is not very difficult to see that $\mathrm{ind}_{\mathbb{Z}_2}(Y) = \mathrm{ind}_{\mathbb{Z}_2}((\mathbb{R}^d)^2_\Delta) \leq d-1$. Indeed, there is a simple $\mathbb{Z}_2$-map $g\colon (\mathbb{R}^d)^2_\Delta \to S^{d-1}$ given by $(\boldsymbol{x}_1, \boldsymbol{x}_2) \mapsto \frac{\boldsymbol{x}_1 - \boldsymbol{x}_2}{\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|}$.

The space $X = (\sigma^{d+1})^2_\Delta$ is more complicated. If we work out the structure of X for $d = 1$, we get a hexagon, i.e. an S^1:

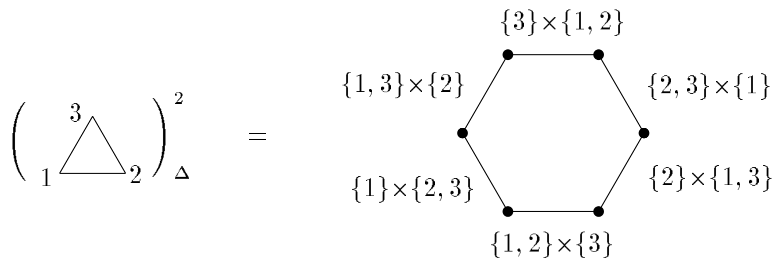

With some more effort, one can find out that for $d = 2$, X can be represented as the boundary of a nice 3-dimensional polytope, as is sketched below:

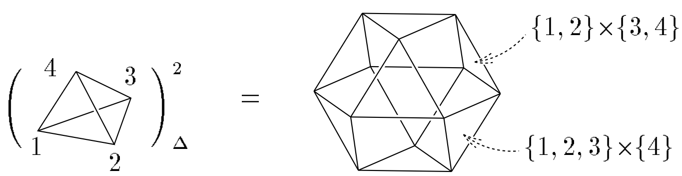

So in this case $X \cong S^2$. In general, one can prove geometrically that $X \cong S^d$ for all d (Exercise 5). This is good, since S^d is $(d-1)$-connected and therefore $\mathrm{ind}_{\mathbb{Z}_2}(X) \geq d$ by Proposition 5.2.2(iv).

As the above 3-dimensional picture for $d = 2$ indicates, the structure of the deleted product $(\sigma^{d+1})^2_\Delta$ is not very simple. In more complicated cases, the deleted products would be even harder to handle. Moreover, in some applications they are not sufficiently connected. In the next section, we introduce another construction, the deleted join, which looks less natural but has significant advantages over the deleted product. We then redo the proof of the topological Radon theorem using deleted joins.

Notes. Surveys on Radon's theorem and its relatives are Eckhoff [Eck79] and [Eck93].

The original proof of the topological Radon theorem by Bajmóczy & Bárány [BB79] is different from the one shown above. They construct a continuous map $g\colon S^d \to \|\sigma^{d+1}\|$ such that for every $\boldsymbol{x} \in S^d$, $\mathrm{supp}(g(\boldsymbol{x})) \cap \mathrm{supp}(g(-\boldsymbol{x})) = \emptyset$, and then they apply the Borsuk-Ulam theorem to $f \circ g\colon S^d \to \mathbb{R}^d$.

Exercises

1. (a) Prove that $\mathrm{ind}_{\mathbb{Z}_2}((\mathbb{R}^d)^2_\Delta) \geq d-1$.

   (b) Check that S^{d-1} is a deformation retract of $(\mathbb{R}^d)^2_\Delta$.

2. Enumerate all the possible configurations for Radon's theorem in dimensions $d = 3$ and $d = 4$.

3. Let $A \subseteq \mathbb{R}^n$ be a set and let $f\colon \mathbb{R}^n \to \mathbb{R}^m$ be an affine map. Show that $\mathrm{conv}(f(A)) = f(\mathrm{conv}(A))$.

4. Let P and Q be convex polytopes in $\mathbb{R}^d$, and let $P + Q = \{\boldsymbol{x} + \boldsymbol{y} : \boldsymbol{x} \in P, \boldsymbol{y} \in Q\}$ be their *Minkowski sum.*

(a) Prove that $P+Q$ is a convex polytope.

(b) Prove that each face of $P+Q$ is of the form $F+G$, where F is a face of P and G is a face of Q.

5. Let $S := \|\sigma^d\| \subset \mathbb{R}^d$ be a (geometric) d-dimensional simplex, and let $P := S+(-S) = \{\boldsymbol{x}-\boldsymbol{y} : \boldsymbol{x}, \boldsymbol{y} \in S\}$.

(a) Verify that P is a d-dimensional convex polytope.

(b) Show that each point $\boldsymbol{x} \in \partial P$ has a *unique* representation in the form $\boldsymbol{x} = \boldsymbol{x}_1 - \boldsymbol{x}_2$, where $\boldsymbol{x}_1, \boldsymbol{x}_2 \in S$ satisfy $\operatorname{supp}(\boldsymbol{x}_1) \cap \operatorname{supp}(\boldsymbol{x}_2) = \emptyset$.

(c) Prove that the deleted product $\|(\sigma^d)^2_\Delta\|$ is homeomorphic to ∂P and, consequently, to S^{d-1}.

5.4 Deleted joins

In this section we introduce deleted joins of simplicial complexes and of spaces and we give another, simpler, proof of the topological Radon theorem 5.3.2. While in the proof using deleted products we have delegated a nontrivial geometric part to the exercises, here we give a full proof.

We begin with the deleted join of a simplicial complex, which is the simplicial complex consisting of the joins of all ordered pairs of disjoint simplices:

5.4.1 Definition ((Twofold) deleted join of a simplicial complex). *Let K be a simplicial complex. The (twofold) deleted join of K is*

$$\mathsf{K}^{*2}_\Delta := \{F_1 {*} F_2 : F_1, F_2 \in \mathsf{K}, F_1 \cap F_2 = \emptyset\} \subseteq K^{*2}.$$

(Recall our convention that $F_1 {*} F_2$ denotes the disjoint union of F_1 and F_2; F_1 comes from the first copy of K and F_2 from the second copy.) The polyhedron of K^{*2}_Δ can be written

$$\|\mathsf{K}^{*2}_\Delta\| = \{t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2 : \boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathsf{K}, \operatorname{supp}(\boldsymbol{x}_1) \cap \operatorname{supp}(\boldsymbol{x}_2) = \emptyset, t \in [0,1]\}.$$

The $\mathbb{Z}_2$-action ν given by the exchange of coordinates, $\nu\colon t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2 \mapsto (1-t)\boldsymbol{x}_2 + t\boldsymbol{x}_1$, makes K^{*2}_Δ into a **free simplicial $\mathbb{Z}_2$-complex**.

Let us have a few examples.

- The deleted join $(\sigma^0)^{*2}_\Delta$ of a single point (the 0-dimensional simplex) consists of two disjoint points.
- The deleted join $(S^0)^{*2}_\Delta$ of two points (a 0-sphere $S^0 = \{\emptyset, \{1\}, \{2\}\}$) is a disjoint union of two edges. In fact, this can be seen from the next figure, which shows, from left to right, the disjoint union of two copies of S^0 (four points), their join (a circle consisting of four edges), and the deleted join.

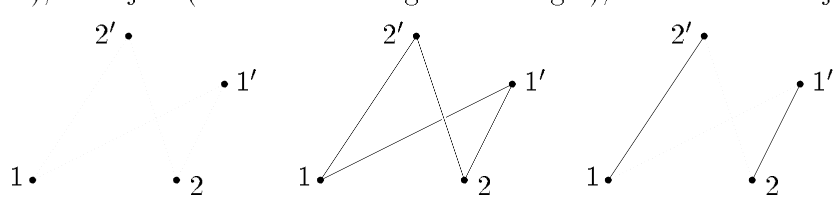

The maximal simplices are $\{1\}*\{2'\}$ and $\{2\}*\{1'\}$. The $\mathbb{Z}_2$-action ν exchanges them.

- The deleted join $(\sigma^1)^{*2}_\Delta$ of an edge is the perimeter of a square. To illustrate this, our figure below shows, from left to right, the disjoint union of two edges, their join (a solid tetrahedron), and the deleted join (as a subcomplex of the tetrahedron).

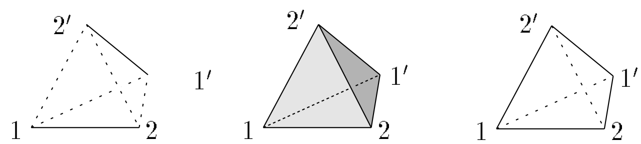

The maximal (1-dimensional) simplices are $\varnothing*\{1',2'\}$, $\{1,2\}*\varnothing$, $\{1\}*\{2'\}$ and $\{2\}*\{1'\}$, where 1 and 2 denote the vertices of σ^1. The $\mathbb{Z}_2$-action ν is the symmetry around the center of the square.

In the proof of the topological Radon's theorem, we will need to compute the deleted join of a simplex. Unlike the deleted product, this is very easy.

5.4.2 Lemma. *Let K and L be simplicial complexes. We have*

$$(\mathsf{K} * \mathsf{L})^{*2}_\Delta = \mathsf{K}^{*2}_\Delta * \mathsf{L}^{*2}_\Delta.$$

Proof. Clear from the definition.

5.4.3 Corollary. $\|(\sigma^n)^{*2}_\Delta\| \cong S^n$.

Proof. We have $\sigma^n = (\sigma^0)^{*(n+1)}$. By Lemma 5.4.2 we obtain

$$((\sigma^0)^{*(n+1)})^{*2}_\Delta = ((\sigma^0)^{*2}_\Delta)^{*(n+1)} = (S^0)^{*(n+1)} \cong S^n.$$

The last homeomorphism is the homeomorphism of the boundary of the cross-polytope in $\mathbb{R}^{n+1}$ with the n-sphere.

We will also need the deleted join of a space.

5.4.4 Definition ((Twofold) deleted join of a space). *Let Z be a topological space. The (twofold) deleted join of X is*

$$X^{*2}_\Delta := X^{*2} \setminus \{\tfrac{1}{2}x + \tfrac{1}{2}x : x \in X\}.$$

A **free $\mathbb{Z}_2$-action** *ν on X^{*2}_Δ is given by $\nu\colon tx_1 + (1-t)x_2 \mapsto (1-t)x_2 + tx_1$.*

Warning. Note the distinction between the deleted join of a simplicial complex and of a space (we had a similar distinction for deleted products). We have $\|\mathsf{K}^{*2}_\Delta\| \subseteq \|\mathsf{K}\|^{*2}_\Delta$, but the inclusion is proper (except for trivial cases)! (On the other hand, these two spaces are homotopy equivalent and have the same $\mathbb{Z}_2$-index; see Exercise 1.) Perhaps one should distinguish these two notions by different notation, but this would add further symbols to learn. Moreover, we need the deleted join of a space exclusively for the case $X = \mathbb{R}^d$, and so no confusion should arise; actually we only need to bound $\mathrm{ind}_{\mathbb{Z}_2}((\mathbb{R}^d)^{*2}_\Delta)$ from above.

5.4.5 Lemma (Deleted join of $\mathbb{R}^d$). *There is a $\mathbb{Z}_2$-map $g\colon (\mathbb{R}^d)^{*2}_\Delta \to S^d$, and consequently $\mathrm{ind}_{\mathbb{Z}_2}((\mathbb{R}^d)^{*2}_\Delta) \le d$. (It can actually be shown that the index equals d.)*

Proof. There are several ways of doing this. We exhibit a $\mathbb{Z}_2$-map $h\colon (\mathbb{R}^d)^{*2}_\Delta \to (\mathbb{R}^{d+1})^2_\Delta$; a $\mathbb{Z}_2$-map $(\mathbb{R}^{d+1})^2_\Delta \to S^d$ was shown in the previous section.

We recall from Proposition 4.2.4 that the join $X * Y$ can be represented geometrically if X and Y are placed into some $\mathbb{R}^n$ as bounded subsets of two skew affine subspaces U and V. In our case, $\mathbb{R}^d$ is unbounded, and so we first map it homeomorphically into the ball B^d and consider the (larger) deleted join $(B^d)^{*2}_\Delta$ instead.

For the geometric representation, we need two skew d-dimensional subspaces, and to preserve the $\mathbb{Z}_2$ symmetry, we choose them in $\mathbb{R}^{2d+2} = (\mathbb{R}^{d+1})^2$. Namely, we define the mappings $\psi_1, \psi_2\colon \mathbb{R}^d \to \mathbb{R}^{2d+2}$ by

$$\psi_1(\boldsymbol{x}) := (1, x_1, \ldots, x_d, 0, 0, \ldots, 0), \quad \psi_2(\boldsymbol{y}) := (0, 0, \ldots, 0, 1, y_1, \ldots, y_d).$$

Then $U_1 := \psi_1(\mathbb{R}^d)$ and $U_2 := \psi_2(\mathbb{R}^d)$ are d-dimensional skew subspaces, and we can insert the two copies of B^d into them: $X_i := \psi_i(B^d)$, $i = 1, 2$. We define $h\colon (B^d)^{*2}_\Delta \to (\mathbb{R}^{d+1})^2$ by $h\colon t\boldsymbol{x} + (1-t)\boldsymbol{y} \mapsto t\psi_1(\boldsymbol{x}) + (1-t)\psi_2(\boldsymbol{y})$. This mapping is continuous by Proposition 4.2.4, is obviously a $\mathbb{Z}_2$-map, and goes into $(\mathbb{R}^{d+1})^2_\Delta$ since the equality $(t, tx_1, \ldots, tx_d) = (1-t, (1-t)y_1, \ldots, (1-t)y_d)$ implies $t = \frac{1}{2}$ and $\boldsymbol{x} = \boldsymbol{y}$, which are exactly the points removed from the deleted join.

Proof of the topological Radon theorem. Now we have everything ready. As before, we assume for contradiction that there is a continuous map $f\colon \|\sigma^{d+1}\| \to \mathbb{R}^d$ where the images of vertex-disjoint faces never intersect. Instead of passing to the mapping f_2 of Cartesian products, we now pass to the mapping of joins:

$$\begin{aligned} f^{*2} := f * f\colon\ \|\sigma^{d+1}\|^{*2} &\to (\mathbb{R}^d)^{*2} \\ t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2 &\mapsto t f(\boldsymbol{x}_1) + (1-t) f(\boldsymbol{x}_2). \end{aligned}$$

If we restrict f^{*2} to the deleted join $X := \|(\sigma^{d+1})^{*2}_\Delta\|$, the image surely contains no point of the form $\frac{1}{2}\boldsymbol{y} + \frac{1}{2}\boldsymbol{y}$ (since f never sends points $\boldsymbol{x}_1, \boldsymbol{x}_2$ from disjoint faces to the same point $\boldsymbol{y} \in \mathbb{R}^d$). Therefore, f^{*2}_Δ can be regarded as a $\mathbb{Z}_2$-map $X \to Y$, where $Y := (\mathbb{R}^d)^{*2}_\Delta$.

We have computed the indices in advance: $\mathrm{ind}_{\mathbb{Z}_2}(X) = \mathrm{ind}_{\mathbb{Z}_2}((\sigma^{d+1})^{*2}_\Delta) = d{+}1$ (Corollary 5.4.3) and $\mathrm{ind}_{\mathbb{Z}_2}(Y) \le d$ by Lemma 5.4.5. Hence $X \overset{\mathbb{Z}_2}{\not\to} Y$ and we have a contradiction proving the topological Radon theorem.

Exercises

1. (Deleted join of a simplicial complex and of its polyhedron)
   (a) Construct a $\mathbb{Z}_2$-map of $\|\sigma^n\|^{*2}_\Delta \to \|(\sigma^n)^{*2}_\Delta\|$; proceed by induction on n.
   (b) Let K be a finite simplicial complex. Show that $\|\mathsf{K}\|^{*2}_\Delta \xrightarrow{\mathbb{Z}_2} \|\mathsf{K}^{*2}_\Delta\|$.

(c) Show that the spaces in (b) are homotopy equivalent; namely, a suitable $\mathbb{Z}_2$-map as in (b) is a homotopy inverse to the obvious insertion $\|\mathsf{K}^{*2}_{\Delta}\| \to \|\mathsf{K}\|^{*2}_{\Delta}$.

5.5 The Van Kampen–Flores theorem

In the proof of the topological Radon theorem in the previous section, we showed that if $\mathrm{ind}_{\mathbb{Z}_2}((\sigma^{d+1})^{*2}_{\Delta}) > \mathrm{ind}_{\mathbb{Z}_2}((\mathbb{R}^d)^{*2}_{\Delta}) = d$, then every continuous map $\|\sigma^{d+1}\| \to \mathbb{R}^d$ identifies two points with disjoint supports. This part of the proof works for any simplicial complex K in place of σ^{d+1}, and gives

5.5.1 Proposition (Nonembeddability and index of the deleted join). *Let K be a simplicial complex. If $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) > d$, then for every continuous mapping $f\colon \|K\| \to \mathbb{R}^d$, the images of some two disjoint faces of K intersect. In particular, $\mathbb{R}^d$ contains no subspace homeomorphic to $\|\mathsf{K}\|$.*

To apply this proposition, one must bound above the index of the deleted join. We begin with an important class of examples.

5.5.2 Theorem (The Van Kampen–Flores theorem; [vK32] [Flo34]). *For all $d \geq 1$, the simplicial complex $\mathsf{K} := (\sigma^{2d+2})^{\leq d}$, i. e. the d-skeleton of the $(2d+2)$-dimensional simplex, cannot be embedded into $\mathbb{R}^{2d}$. In fact, for any continuous map $f\colon \|\mathsf{K}\| \to \mathbb{R}^{2d}$, the images of some two disjoint faces of K intersect.*

In Theorem 1.6.1, we have embedded an arbitrary d-dimensional simplicial complex into $\mathbb{R}^{2d+1}$, and the Van Kampen–Flores theorem shows that this dimension is best possible in general. The case $d = 1$ tells us that in any drawing of the complete graph K_5 in the plane, some two vertex-disjoint edges intersect.

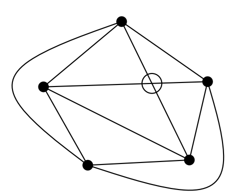

This is well-known and can be proved in an elementary way (without the Jordan curve theorem; see, for example, [Tho92]), although the proofs usually offered in graph theory courses are not rigorous.

We present two proofs of the Van Kampen–Flores theorem. In both of them, we show $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) > 2d$ and apply Proposition 5.5.1. The first proof, in this section, analyzes K^{*2}_{Δ} in detail and shows that it is actually homeomorphic to S^{2d+1}. On the way, we introduce an interesting class of triangulations of S^{2d+1}. This part is optional and is not needed for the further development. The second proof, given in the next section, contains important ideas which will be used in several other applications.

The Bier spheres. This first proof of the Van Kampen–Flores theorem is similar to the original proof of Flores. To analyze the complex K^{*2}_{Δ}, we consider a more general construction, due to Bier [Bie], which associates an $(n-2)$-dimensional triangulated sphere on $2n$ vertices with every simplicial complex on n vertices. It is simple but ingenious, and we will also present another application of it.

Recall that $2^{[n]}$ denotes the system of all subsets of $[n] = \{1, 2, \ldots, n\}$. A simplicial complex with vertex set $[n]$ is a nonempty subset $\mathsf{K} \subseteq 2^{[n]}$. Strictly speaking, a vertex of such K is not an element $i \in [n]$ but rather the 0-dimensional simplex $\{i\}$. Up until now, there was no need to distinguish this, since we always tacitly assumed that all elements of the ground set are 0-dimensional simplices. But now it does make some difference, since although we allow that $\{i\} \notin \mathsf{K}$ for some $i \in [n]$, we still want to speak of simplicial complexes with the ground set $[n]$. In order to make the formulas shorter, let us write $\overline{F}$ for $[n]\setminus F$, where $F \subseteq [n]$.

5.5.3 Definition. *Let $\mathsf{K} \subset 2^{[n]}$ be a simplicial complex on the ground set $[n]$. The* **Alexander dual** *of K is the simplicial complex $B(\mathsf{K}) \subseteq 2^{[n]}$ that consists of the complements of the non-simplices of K:*

$$B(\mathsf{K}) \;:=\; \{G \subseteq [n] : \overline{G} \notin \mathsf{K}\} \;=\; \{\overline{H} : H \in 2^{[n]}\setminus\mathsf{K}\}.$$

The **Bier sphere** *associated with K is defined as the deleted join*

$$\begin{aligned}\mathrm{Bier}_n(\mathsf{K}) \;:=\; \mathsf{K} *_{\Delta} B(\mathsf{K}) &= \{F{*}G : F \in \mathsf{K}, \overline{G} \notin \mathsf{K},\ F \cap G = \emptyset\} \\ &= \{F{*}\overline{H} : F \in \mathsf{K},\ H \notin \mathsf{K},\ F \subset H\}.\end{aligned}$$

In this construction, neither K nor $B(\mathsf{K})$ has to have all elements $i \in [n]$ as vertices—we just assume that their vertex sets are contained in $[n]$. However, if i is not a vertex of K (that is, $\{i\} \notin \mathsf{K}$), then $[n]\setminus\{j\}$ is never a face of K for $j \neq i$, and hence j is a vertex of $B(\mathsf{K})$. It follows easily that $\mathrm{Bier}_n(\mathsf{K})$ is a simplicial complex with at least n vertices. You may also note that here we form a deleted join of complexes that are different and may, in general, even have distinct vertex sets, but they have the same ground set $[n]$.

5.5.4 Theorem (Bier sphere is a sphere). *For every simplicial complex $\mathsf{K} \subset 2^{[n]}$, the simplicial complex $\mathrm{Bier}_n(\mathsf{K})$ is an $(n-2)$-sphere with at most $2n$ vertices.*

Before proving this theorem, let us check that it yields what we want for the proof of the Van Kampen–Flores theorem.

5.5.5 Example (The Flores sphere). Take $n = 2d{+}3$ and $\mathsf{K} = \binom{[n]}{<d+1}$, the d-skeleton of the $(2d{+}2)$-dimensional simplex. In this case $B(\mathsf{K}) = \mathsf{K}$, and hence by Theorem 5.5.4, $\mathrm{Bier}_n(\mathsf{K}) = \mathsf{K}^{*2}_{\Delta}$ is a $(2d{+}1)$-sphere.

For example, for $d = 0$ and $n = 3$, we have $\mathsf{K} = [3]$ (three disjoint points), and the deleted join K^{*2}_{Δ} is a hexagon:

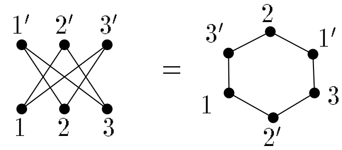

5.5.6 Lemma. *The facets (maximal simplices) of the Bier sphere* $\mathrm{Bier}_n(\mathsf{K})$ *are*

$$F{*}\overline{H} \quad \text{where} \quad F \subset H,\ F \in \mathsf{K},\ H \notin \mathsf{K}, \quad \text{and} \quad |H\backslash F| = 1.$$

In particular, $\mathrm{Bier}_n(\mathsf{K})$ *is a pure complex of dimension* $n{-}2$ *(i. e. each simplex is contained in a maximal* $(n{-}2)$*-dimensional simplex).*

Proof. For any face $F_0{*}\overline{H_0} \in \mathrm{Bier}_n(\mathsf{K})$ we can find $F \in \mathsf{K}$ and $H \notin \mathsf{K}$ with

$$F_0 \subseteq F \subset H \subseteq H_0 \qquad \text{and} \qquad |H\backslash F| = 1.$$

We get $F_0{*}\overline{H_0} \subseteq F{*}\overline{H} \in \mathrm{Bier}_n(\mathsf{K})$. The size of the face $F{*}\overline{H}$, namely $|F \cup \overline{H}| = |F| + n - |H| = n{-}1$, is clearly maximal, since $F \subset H$ and hence $|H\backslash F| \geq 1$.

5.5.7 Examples. The simplest complex to study is probably the empty one: $\mathsf{K} = \{\varnothing\}$. For this we get $\mathrm{Bier}_n(\mathsf{K}) = B(\mathsf{K}) = 2^{[n]}\backslash[n]$, the boundary complex of an $(n{-}1)$-dimensional simplex, with n vertices. Thus $\|\mathrm{Bier}_n(\{\varnothing\})\| \cong S^{n-2}$.

If we take $\mathsf{K} = 2^{[n-1]} = \sigma^{n-2}$, then $B(\mathsf{K}) = \mathsf{K}$, and thus $\mathrm{Bier}_n(\mathsf{K}) = (\sigma^{n-2})^{*2}_{\Delta}$ is the deleted join of an $(n{-}2)$-simplex. This is the simplicial sphere given by the boundary of an $(n{-}1)$-dimensional crosspolytope, with $2(n{-}1)$ vertices, by Corollary 5.4.3.

Proof of Theorem 5.5.4. Let $F \notin \mathsf{K}$ be any inclusion-minimal non-face of K. Then $\mathsf{K} \cup \{F\}$ is a simplicial complex as well, and for the maximal (i. e. $(n{-}2)$-dimensional) simplices of the Bier spheres we find

$$\begin{aligned} \mathrm{Bier}_n(\mathsf{K} \cup \{F\})^{n-2} \quad = \quad & \mathrm{Bier}_n(\mathsf{K})^{n-2} \ \backslash \ \{(F\backslash\{i\}) * \overline{F} : i \in F\} \\ & \qquad \cup \ \{F * \overline{F \cup \{j\}} : j \notin F\}. \end{aligned}$$

The vertex sets of the simplices affected by this operation (added or removed) are all contained in $V_F = \{\{i\}{*}\varnothing : i \in F\} \cup \{\varnothing{*}\{j\} : j \in \overline{F}\}$. The subcomplex L_1 of $\mathrm{Bier}_n(\mathsf{K})$ induced by the vertex set V_F is

$$\mathsf{L}_1 = (2^F\backslash\{F\}) * 2^{\overline{F}}$$

while the corresponding subcomplex in $\mathrm{Bier}_n(\mathsf{K} \cup \{F\})$ is

$$\mathsf{L}_2 = 2^F * (2^{\overline{F}}\backslash\{\overline{F}\}).$$

Their common part is

$$\mathsf{L}_0 = \mathsf{L}_1 \cap \mathsf{L}_2 = (2^F\backslash\{F\}) * (2^{\overline{F}}\backslash\{\overline{F}\}).$$

This is the join of the boundary of a $(k{-}1)$-simplex, $k = |F|$, with the boundary of an $(n{-}k{-}1)$-simplex, so $\|\mathsf{L}_0\| \cong S^{n-3}$. Both L_1 and L_2 are triangulations of an $(n{-}2)$-ball bounded by this S^{n-3}. For example, for $n = 4$ and $F = \{1, 2\}$, the geometric picture in $\mathbb{R}^3$ is

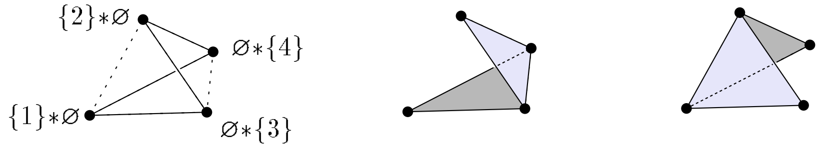

L_0 L_1 L_2

and another possibility, with $F = \{1\}$, is

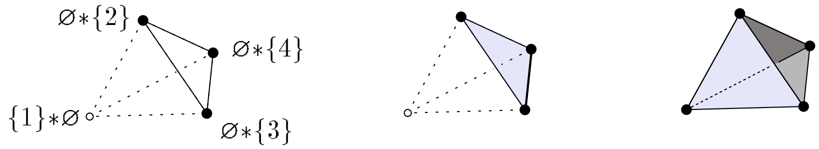

L_0 L_1 L_2

Further we note that a simplex having a vertex outside of V_F never contains a simplex in $\mathsf{L}_1 \setminus \mathsf{L}_0$ (or in $\mathsf{L}_2 \setminus \mathsf{L}_0$). So both $\|\mathsf{L}_1\|$ and $\|\mathsf{L}_2\|$ are $(n-2)$-balls glued to the rest of the Bier sphere by the $(n-3)$-sphere $\|\mathsf{L}_0\|$, and $\|\mathrm{Bier}_n(\mathsf{K})\|$ and $\|\mathrm{Bier}_n(\mathsf{K} \cup \{F\})\|$ are homeomorphic.

The re-triangulation of the ball bounded by the sphere $\|\mathsf{L}_0\|$ is called a *bistellar* operation. It can be geometrically interpreted in $\mathbb{R}^{n-2}$: consider a $(k-1)$-simplex A_1, $k = |F|$, and an $(n-k-1)$-simplex A_2, placed in $\mathbb{R}^{n-2}$ so that they intersect at a single point belonging to their relative interiors (this is a "Radon configuration" as in Theorem 5.3.1). The bistellar operation corresponds to switching between two triangulations of $\mathrm{conv}(A_1 \cup A_2)$. This convex polytope is a projection of an $(n-1)$-simplex in $\mathbb{R}^{n-1}$ into $\mathbb{R}^{n-2}$, and the triangulations correspond to the "top" and "bottom" views of that simplex. For $n = 4$, the possible operations are: switching the diagonal in a quadrilateral ($k = 2$), a stellar subdivision of a triangle (adding a new vertex: $k = 1$), and its inverse operation, namely removing such a stellar subdivision (and thus deleting one vertex: $k = 3$).

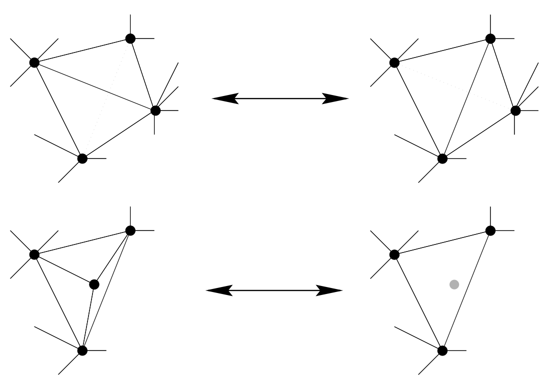

Note how this corresponds to the two 3-dimensional pictures above.

Since every simplicial complex $\mathsf{K} \subset 2^{[n]}$ can be generated from $\{\varnothing\}$ by adding minimal non-faces, we see that $\mathrm{Bier}_n(\mathsf{K})$ is homeomorphic to $\mathrm{Bier}_n(\{\varnothing\})$. This is an $(n-2)$-sphere by Example 5.5.7.

Many PL-spheres. This proof establishes more. Bistellar operations preserve the property of spheres to be *piecewise-linear* (PL); that is, to have a subdivision that is also a subdivision of a simplex boundary. In fact, by results of Pachner [Pac86], a sphere is PL *if and only if* it can be generated from the boundary of a simplex (such as $\mathrm{Bier}_n(\{\varnothing\})$) using only bistellar operations.

We have produced *many* $(n-2)$-spheres with $2n$ vertices. Let $n \geq 4$ and consider all simplicial complexes $\mathsf{K} \subset 2^{[n]}$ that contain all faces of dimension $\lfloor n/2 \rfloor - 2$ and some of the $(\lfloor n/2 \rfloor - 1)$-dimensional faces:

$$\binom{[n]}{\leq \lfloor n/2 \rfloor - 1} \subseteq \mathsf{K} \subseteq \binom{[n]}{\leq \lfloor n/2 \rfloor}.$$

The number of such K is

$$2^{\binom{n}{\lfloor n/2 \rfloor}} > 2^{2^n/n}.$$

The Bier spheres $\mathrm{Bier}_n(\mathsf{K})$ have exactly $2n$ vertices. For every Bier sphere $\mathrm{Bier}_n(\mathsf{K})$, there are not more than $(2n)!/n! < (2n)^n < 2^{n^2}$ different complexes K that yield an isomorphic Bier sphere, and hence there are more than

$$2^{4^n/n - n^2}$$

non-isomorphic simplicial $(n-2)$-spheres with $2n$ vertices, a doubly exponential number! This shows that most of the simplicial $(n-2)$-spheres on $2n$ vertices cannot be realized as boundary complexes of $(n-1)$-dimensional convex polytopes; they cannot be made "straight." This is because the number of different combinatorial types of $(n-1)$-dimensional polytopes with $2n$ vertices is no larger than

$$2^{4n^3}.$$

Such bound can be derived from the results of Oleinik and Petrovskiĭ, Milnor, and Thom on the topological complexity of algebraic varieties; see Goodman & Pollack [GP86, last line of p. 222].

Notes. The Van Kampen–Flores theorem was proved by Van Kampen [vK32] and Flores [Flo34] independently, at the same time. An exposition of Flores' proof can be found in Grünbaum's book on convex polytopes [Grü67, Sect. 11.2], while our development in the next section can be traced back to the Van Kampen's proof.

Realizability of simplicial complexes in $\mathbb{R}^d$ is a very interesting and largely unexplored area. For $d = 2$, we have the well-developed theory of planar graphs and of various measures of non-planarity of a graph (the crossing number etc.), but even higher-dimensional analogues of very basic theorems about planar graphs remain unclear. The behavior in higher dimensions can also be very different from the planar case.

For example, as is well-known, any planar graph has a planar drawing where all edges are straight segments. While every d-dimensional simplicial complex embeds into $\mathbb{R}^{2d+1}$, and even linearly, Brehm and Sarkaria [BS92] proved that for each $d \geq 2$ and $r \geq 1$, there are d-dimensional simplicial complexes K realizable in $\mathbb{R}^{2d}$ but such that the rth barycentric subdivision $\mathrm{sd}(\mathrm{sd}(\ldots\mathrm{sd}(\mathsf{K})\ldots))$ cannot be realized linearly in $\mathbb{R}^{2d}$ (i. e. so that the embedding is affine on each simplex). So a piecewise linear realization in $\mathbb{R}^{2d}$ must have very many pieces. For

$d = 2^k$, there are even d-dimensional triangulations of manifolds with boundary that embed in $\mathbb{R}^{2d-1}$ but not linearly (while all triangulations of d-manifolds do embed linearly in $\mathbb{R}^{2d}$).

Necessary and sufficient topological conditions for realizability of a d-dimensional simplicial complex in $\mathbb{R}^{2d}$, $d \geq 3$, were stated by Van Kampen in 1932 and by Flores in 1933, and proved in detail independently by Shapiro [Sha57] and by Wu [Wu65]. (The case $d = 2$ is really exceptional.) Interesting necessary conditions for *linear* realizability of simplicial complexes were found by Novik [Nov00].

A planar graph on n vertices has at most $3n-6$ edges. Is it true that any simplicial complex on n vertices realizable in $\mathbb{R}^d$ has at most $C_d n^{\lceil d/2 \rceil}$ simplices, for some C_d depending on d but not on n? If true, this would be best possible, as is witnessed by the boundary complex of a cyclic $(d+1)$-polytope with one d-simplex removed, but the problem remains open. For $d = 3$, there is an elementary proof (Dey and Edelsbrunner [DE94]): assume that the embedding is piecewise linear, say, and consider a tiny sphere around each vertex v; then the intersections of the edges and triangles adjacent to v give a planar graph drawn on the sphere, with $O(n)$ edges. A study of some embedding questions for higher-dimensional complexes by elementary methods is Dey and Pach [DP98].

The first construction of "many" simplicial spheres was given by Kalai [Kal88]: for that, Kalai used the cyclic polytopes in the place where here we (combinatorially) deal with the crosspolytopes.

5.6 Sarkaria's inequality

In this section we give another proof of the Van Kampen–Flores theorem 5.5.2, in which we demonstrate a powerful trick for bounding the $\mathbb{Z}_2$-index.

Recall that we need to prove $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) > 2d$, where K is the d-skeleton of the $(2d+2)$-simplex. Our simplicial complex $\mathsf{L}_0 := \mathsf{K}^{*2}_{\Delta}$ is a subcomplex of $\mathsf{L} := (\sigma^{2d+2})^{*2}_{\Delta}$. We already know that $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{L}) = 2d+2$ (Corollary 5.4.3). The idea is to look at the complement of L_0 within L, see that it is "small," and conclude that L_0 must be "large." One immediate problem is that the complement $\mathsf{L} \setminus \mathsf{L}_0$ is not a simplicial complex. Yet it can be used to define a simplicial complex, namely the order complex of the partially ordered set $(\mathsf{L} \setminus \mathsf{L}_0, \subseteq)$.

In the following lemma, we consider a slightly more symmetric situation, where the simplices of L are covered by two arbitrary subsets. For an arbitrary family $\mathcal{F}$ of finite sets, let $\Delta_0(\mathcal{F})$ with $\subseteq$ denote the order complex of the poset $(\mathcal{F} \setminus \{\varnothing\}, \subseteq)$. If it is clear that $\varnothing \notin \mathcal{F}$, we write just $\Delta(\mathcal{F})$.

5.6.1 Lemma. *Let L be a simplicial complex and let $\mathsf{L} = \mathcal{L}_0 \dot\cup \mathcal{L}_1$ be a partition of the simplices of L into two subsets. Then there is a (canonical) simplicial embedding*

$$\iota\colon\ \mathrm{sd}(\mathsf{L}) \longrightarrow \Delta_0(\mathcal{L}_0) * \Delta_0(\mathcal{L}_1).$$

If L is a simplicial $\mathbb{Z}_2$-complex and $\mathcal{L}_0$ and $\mathcal{L}_1$ are both invariant under the $\mathbb{Z}_2$-action, then $\Delta_0(\mathcal{L}_0)$ and $\Delta_0(\mathcal{L}_1)$ are simplicial $\mathbb{Z}_2$-complexes and ι provides a $\mathbb{Z}_2$-map

$$\iota\colon\ \|\mathsf{L}\| \longrightarrow \|\Delta_0(\mathcal{L}_0)\| * \|\Delta_0(\mathcal{L}_1)\|.$$

Let us have a geometric example first. Let L be the 2-simplex and let $\mathsf{L} = \mathcal{L}_0 \dot{\cup} \mathcal{L}_1$ be the partition of its simplices indicated in the picture:

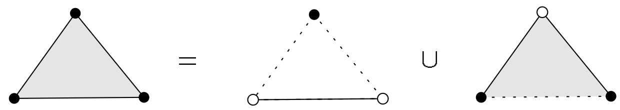

L $\mathcal{L}_0$ $\mathcal{L}_1$

Geometrically, $\Delta_0(\mathcal{L}_0)$ is the subcomplex of the first barycentric subdivision $\mathrm{sd}(\mathsf{L})$ induced by the barycenters of the simplices in $\mathcal{L}_0$. For our examples, we have

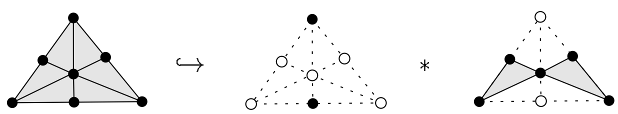

$\mathrm{sd}(\mathsf{L})$ $\Delta_0(\mathcal{L}_0)$ $\Delta(\mathcal{L}_1)$

Note that the vertex sets of $\Delta_0(\mathcal{L}_0)$ and of $\Delta_0(\mathcal{L}_1)$ form a partition of the vertex set of $\mathrm{sd}(\mathsf{L})$; this is just rephrasing of the assumption $\mathsf{L} = \mathcal{L}_0 \dot{\cup} \mathcal{L}_1$.

Proof of Lemma 5.6.1. The vertex set $V(\Delta_0(\mathcal{L}_0) * \Delta_0(\mathcal{L}_1))$ is the union of $V(\Delta_0(\mathcal{L}_0))$ and $V(\Delta_0(\mathcal{L}_1))$ and it equals $V(\mathrm{sd}(\mathsf{L}))$. So, on the level of vertices, we can just set $\iota(F) := F$, $F \in \mathsf{L}$. This map is simplicial: if $\mathcal{C} = \{F_1, F_2, \ldots, F_n\}$ is a chain of simplices of L, $F_1 \subset F_2 \subset \cdots \subset F_n$, it splits into the chains $\mathcal{C}_0 := \mathcal{C} \cap \mathcal{L}_0$ and $\mathcal{C}_1 := \mathcal{C} \cap \mathcal{L}_1$. The concatenation $\mathcal{C}_0 * \mathcal{C}_1$ of these chains is a simplex of the join $\Delta_0(\mathcal{L}_0) * \Delta_0(\mathcal{L}_1)$.

It remains to check the equivariance of ι if L is a simplicial $\mathbb{Z}_2$-complex and $\mathcal{L}_0$, $\mathcal{L}_1$ are invariant subsets of simplices. This is straightforward and is left to the reader.

If we let $\mathcal{L}_0 = \mathsf{L}_0$ be a subcomplex of L, then $\Delta_0(\mathcal{L}_0) = \mathrm{sd}(\mathsf{L}_0)$ is homeomorphic to L_0. Together with Proposition 5.2.2(iii), about the $\mathbb{Z}_2$-index of a join, Lemma 5.6.1 yields

5.6.2 Theorem (Sarkaria's inequality). *Let L be a finite simplicial $\mathbb{Z}_2$-complex and let L_0 be an invariant subcomplex of L. Then we have*

$$\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{L}_0) \geq \mathrm{ind}_{\mathbb{Z}_2}(\mathsf{L}) - \mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathsf{L} \setminus \mathsf{L}_0)) - 1.$$

In combination with Proposition 5.5.1 about non-embeddability and $\mathbb{Z}_2$-index of a deleted join, we obtain

5.6.3 Corollary. *Let* K *be a subcomplex of a simplicial complex* J. *If*

$$\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L})) \leq \mathrm{ind}_{\mathbb{Z}_2}(\mathsf{J}^{*2}_{\Delta}) - d - 2,$$

where $\mathcal{L} := \mathsf{J}^{*2}_{\Delta} \setminus \mathsf{K}^{*2}_{\Delta}$, *then for any continuous map* $f\colon \|\mathsf{K}\| \to \mathbb{R}^d$, *the images of some two disjoint faces of* K *intersect. In particular, for* $\mathsf{J} = \sigma^n$, *we require*

$$\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L})) \leq n - d - 2.$$

Second proof of the Van Kampen–Flores theorem. We use Corollary 5.6.3 for embedding into $\mathbb{R}^{2d}$ with $J := \sigma^{2d+2}$ and K being the d-skeleton of J.

Here the vertices of $\Delta(\mathcal{L})$, which are the simplices of $\mathcal{L}$, have the form $F_1 * F_2$, where $F_1, F_2 \subseteq [2d{+}3]$ are disjoint sets and at least one of them has more than $d{+}1$ vertices. The key observation is that they cannot *both* have more than $d{+}1$ vertices since there is not enough room; the ground set has only $2d{+}3$ points. So the vertices of $\Delta(\mathcal{L})$ naturally fall into two classes: those with $|F_1| \geq d{+}2$ and those with $|F_2| \geq d{+}2$. The $\mathbb{Z}_2$-action on $\Delta(\mathcal{L})$ swaps these two classes.

Let the two vertices of the 0-sphere be 1 and 2. We define a mapping $f\colon V(\Delta(\mathcal{L})) \to S^0$ by

$$f(F_1 * F_2) := \begin{cases} 1 & \text{if } |F_1| \geq d{+}2 \\ 2 & \text{if } |F_2| \geq d{+}2. \end{cases}$$

We claim that f is a simplicial $\mathbb{Z}_2$-map of $\Delta(\mathcal{L})$ to S^0. (This implies that $\Delta(\mathcal{L})$ is disconnected.) It clearly commutes with the $\mathbb{Z}_2$-actions (on S^0, the $\mathbb{Z}_2$-action exchanges 1 and 2). It is simplicial as well, since if $|F_1| \geq d{+}2$ and $F_1' \supseteq F_1$, then $|F_1'| \geq d{+}2$.

Therefore, $\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L})) = 0$ and

$$\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) \geq \mathrm{ind}_{\mathbb{Z}_2}((\sigma^{2d+2})^{*2}_{\Delta}) - \mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L})) - 1 \geq 2d{+}2 - 0 - 1 > 2d.$$

The Van Kampen–Flores theorem is proved.

Notes. The ideas in the proof shown in this section are from Sarkaria's papers [Sar91a], [Sar90]. Our presentation owes much to Živaljević's survey [Živ96], where he isolated "Sarkaria's inequality" and expressed it elegantly using the index.

Exercises

1. Find an example of a simplicial $\mathbb{Z}_2$-complex L and a $\mathbb{Z}_2$-subcomplex L_0 where Sarkaria's inequality 5.6.2 is strict.

5.7 Index, colorings, and another proof of Kneser's conjecture

Corollary 5.6.3 provides a method for showing that a given simplicial complex K cannot be embedded into $\mathbb{R}^d$. To use it, we need to bound above $\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L}))$

with $\mathcal{L} := \mathsf{J}^{*2}_{\Delta} \setminus \mathsf{K}^{*2}_{\Delta}$, where J is a suitable simplicial complex containing K for which we know $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{J}^{*2}_{\Delta})$. Here, surprisingly, we relate $\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathcal{L}))$ to the chromatic number of a certain Kneser graph.

For a set system $\mathcal{S}$, let $\mathrm{MIN}(\mathcal{S})$ denote the system of all sets in $\mathcal{S}$ that are minimal with respect to inclusion (no proper subset is in $\mathcal{S}$). We further recall that $\mathrm{KG}(\mathcal{S})$ denotes the Kneser graph of $\mathcal{S}$, with vertex set $\mathcal{S}$ and with edges connecting disjoint sets.

5.7.1 Lemma. *Let* K *be a subcomplex of a simplicial complex* J *and let* $\mathcal{S} := \mathrm{MIN}(\mathsf{J} \setminus \mathsf{K})$. *Then*

$$\mathrm{ind}_{\mathbb{Z}_2}(\Delta(\mathsf{J}^{*2}_{\Delta} \setminus \mathsf{K}^{*2}_{\Delta})) \leq \chi(\mathrm{KG}(\mathcal{S})) - 1.$$

Proof. Let $m := \chi(\mathrm{KG}(\mathcal{S}))$, and let $c\colon \mathcal{S} \to [m]$ be a proper coloring of $\mathrm{KG}(\mathcal{S})$ with m colors, i.e. $c(F_1) \neq c(F_2)$ whenever $F_1 \cap F_2 = \emptyset$.

Let us write $\mathcal{L} := \mathsf{J}^{*2}_{\Delta} \setminus \mathsf{K}^{*2}_{\Delta}$. We would like to construct a $\mathbb{Z}_2$-map of $\Delta(\mathcal{L})$ into S^{m-1}.

The first trick is to represent the sphere S^{m-1} as the first barycentric subdivision of the deleted join $(\sigma^{m-1})^{*2}_{\Delta}$ (which is correct by Corollary 5.4.3). The required $\mathbb{Z}_2$-map is constructed as a simplicial map

$$g\colon\ \Delta(\mathcal{L}) \ \longrightarrow\ \mathrm{sd}((\sigma^{m-1})^{*2}_{\Delta}).$$

A vertex of the complex on the left-hand side has the form $F_1{*}F_2$, where F_1 and F_2 are disjoint faces of J, at least one of them not belonging to K. A vertex of the complex on the right-hand side is of the form $G_1{*}G_2$, where G_1 and G_2 are disjoint subsets of $[m]$, not both of them empty. We define $g(F_1{*}F_2) := h(F_1){*}h(F_2)$ for a suitable map h assigning subsets of $[m]$ to simplices of J; this automatically guarantees that g is a $\mathbb{Z}_2$-map.

We define

$$h(F) := \{c(G) : G \in \mathcal{S}, G \subseteq F\}.$$

We need to verify that if $F_1{*}F_2 \in \mathcal{L}$, then $h(F_1)$ and $h(F_2)$ are disjoint subsets of $[m]$, not both empty. If $F_1 \cap F_2 = \emptyset$, then $h(F_1) \cap h(F_2) = \emptyset$ as well, for otherwise we would have sets $G_1 \subseteq F_1$ and $G_2 \subseteq F_2$, $G_1, G_2 \in \mathcal{S}$, with $c(G_1) = c(G_2)$ and c would not be a proper coloring of the Kneser graph. The nonemptiness of $h(F_1) \cup h(F_2)$ also follows because we have $h(F) \neq \emptyset$ exactly if $F \in \mathsf{J} \setminus \mathsf{K}$. The map h is monotone with respect to inclusion and so g is a simplicial $\mathbb{Z}_2$-map as claimed.

Putting this together with Corollary 5.6.3, we obtain the following amazing connection between Kneser colorings and embeddability into $\mathbb{R}^d$.

5.7.2 Theorem (Sarkaria's coloring/embedding theorem). *Let* K *be a subcomplex of a simplicial complex* J *and let* $\mathcal{S} := \mathrm{MIN}(\mathsf{J} \setminus \mathsf{K})$. *Then*

$$\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) \geq \mathrm{ind}_{\mathbb{Z}_2}(\mathsf{J}^{*2}_{\Delta}) - \chi(\mathrm{KG}(\mathcal{S})),$$

and consequently, if

$$d \leq \mathrm{ind}_{\mathbb{Z}_2}(\mathsf{J}^{*2}_{\Delta}) - \chi(\mathrm{KG}(\mathcal{S})) - 1,$$

then for any continuous mapping $f\colon \|\mathsf{K}\| \to \mathbb{R}^d$, the images of some two disjoint faces of K *intersect. For* $\mathsf{J} = \sigma^n$*, the condition becomes*

$$d \leq n - \chi(\mathrm{KG}(\mathcal{S})) - 1.$$

Let us first see a few examples of using this remarkable result in proofs of nonrealizability of simplicial complexes.

5.7.3 Example. The Van Kampen–Flores theorem is the special case with $\mathsf{J} = \sigma^{2d+2}$ and $\mathsf{K} = (\sigma^{2d+2})^{\leq d}$. Here $\mathcal{S}$ are all simplices of dimension $d{+}1$, and the Kneser graph $\mathrm{KG}(\mathcal{S})$ has no edges at all, since no two sets in $\mathcal{S}$ are disjoint. So $\chi(\mathrm{KG}(\mathcal{S})) = 1$ and Theorem 5.7.2 gives nonrealizability of K in $\mathbb{R}^{2d}$ as it should.

5.7.4 Example. Let us prove, by this heavy machinery, that the complete bipartite graph $K_{3,3}$ is not planar.

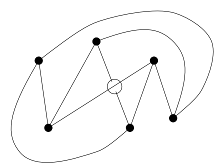

We let the vertex set be $\{1, 2, 3, 1', 2', 3'\}$, J is the 5-simplex on this set, and the maximal simplices of K are $\{i, j'\}$, $i, j = 1, 2, 3$. Then $\mathcal{S}$ consists of the pairs that are not edges of K, i. e. the pairs $\{i, j\}$ or $\{i', j'\}$. We can color the pairs on $\{1, 2, 3\}$ red and the pairs on $\{1', 2', 3'\}$ blue, and so $\chi(\mathrm{KG}(\mathcal{S})) = 2$. Thus $K_{3,3}$ cannot be realized in $\mathbb{R}^d$ for $d \leq 5 - 2 - 1 = 2$.

The index of K^{*2}_{Δ} can easily be computed directly here: observing that $\mathsf{K} = [3]*[3]$, we have $\mathsf{K}^{*2}_{\Delta} = ([3]^{*2}_{\Delta})^{*2}$. Since $[3]^{*2}_{\Delta}$ is a cycle (of length 6), $\|\mathsf{K}^{*2}_{\Delta}\| \cong S^1 * S^1 \cong S^3$.

5.7.5 Example (Non-realizability of $\mathbb{R}\mathbf{P}^2$ in 3-space; Brehm & Sarkaria [BS92]). Let $\mathsf{K} \subseteq 2^{[6]}$ be the pure, 2-dimensional complex whose maximal faces are given by the list

124	125	134	136	156	235	236	246	345	456

This is a remarkable complex. We note four things:

(i) K corresponds to the triangulation of a hexagon drawn below, where opposite vertices and edges on the boundary are identified.

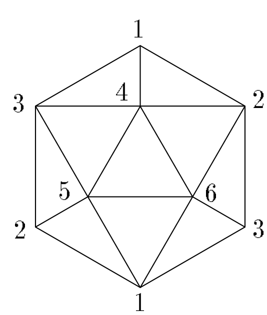

Thus, K triangulates the real projective plane $\mathbb{R}\mathrm{P}^2$. (Another interpretation is that K is the complex obtained by identifying all opposite faces on the boundary of a regular icosahedron. The icosahedron has 12 vertices, 20 edges and 10 triangles, and so the complex K we are looking at has exactly half of these face numbers.)

(ii) K has a complete 1-skeleton: we have $\binom{[6]}{\le 2} \subseteq \mathsf{K}$.

(iii) For triples $F \in \binom{[6]}{3}$ we find that $F \in \mathsf{K}$ if and only if $[6]\backslash F \notin \mathsf{K}$. From this we derive that $B(\mathsf{K}) = \mathsf{K}$, and thus $\mathrm{Bier}_6(\mathsf{K}) = \mathsf{K}^{*2}_{\Delta}$. Therefore, $\mathrm{ind}(\mathsf{K}^{*2}_{\Delta}) = 4$, and Proposition 5.5.1 gives non-realizability in $\mathbb{R}^3$.

(iv) The system $\mathcal{S}$ of minimal non-faces is $\binom{[6]}{3}\backslash\mathsf{K}$, and the Kneser graph is again trivial, since $\mathcal{S}$ has no disjoint simplices. Thus, from Theorem 5.7.2 we obtain another proof of non-realizability of $\|\mathsf{K}\|$ in $\mathbb{R}^3$.

Since embeddability is independent of the triangulation, we have proved that the real projective plane $\mathbb{R}\mathrm{P}^2$ has no embedding into $\mathbb{R}^3$.

Third proof of Kneser's conjecture. Sarkaria's theorem can be used not only for proving the impossibility of an embedding from the existence of a Kneser coloring, but also the other way round.

Let $\mathcal{S} := \binom{[n]}{k}$ be given, and choose the simplicial complex K accordingly to consist of all proper subsets of the sets in $\mathcal{S}$. That is, K is the $(k-2)$-skeleton of σ^{n-1}, and in particular, $\dim(\mathsf{K}) = k-2$. By the geometric realization theorem 1.6.1, $\|\mathsf{K}\|$ can be realized in $\mathbb{R}^{2(k-2)+1} = \mathbb{R}^{2k-3}$. Theorem 5.7.2 gives $\chi(\mathrm{KG}(\mathcal{S})) \ge n-2k+2$, as it should be.

Alternatively, we can avoid speaking about an embedding and use the first inequality in Theorem 5.7.2 directly. It gives $\chi(\mathrm{KG}(\mathcal{S})) \ge \mathrm{ind}_{\mathbb{Z}_2}((\sigma^{n-1})^{*2}_{\Delta}) - \mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) = n-1-\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) \ge n-1-\dim(\mathsf{K}^{*2}_{\Delta}) = n-2k+2$.

Notes. The main results of this section are due to Sarkaria [Sar91a] and [Sar90] (who formulated them for concrete examples rather than as general statements).

Similar to Example 5.7.5, one can prove that the complex projective plane $\mathbb{C}\mathrm{P}^2$ cannot be realized in $\mathbb{R}^6$. There is a 9-vertex triangulation $\mathsf{K} \subseteq 2^{[9]}$ of $\mathbb{C}\mathrm{P}^2$. It is a pure, 4-dimensional simplicial complex with 9 vertices such that $B(\mathsf{K}) = \mathsf{K}$; consequently, the deleted join is a Bier sphere. Again we get that there are no disjoint non-faces, which implies that there is no embedding of this complex, and thus of $\mathbb{C}\mathrm{P}^2$, into $\mathbb{R}^d$ for $d \le 8-1-1 = 6$. (See Kühnel & Banchoff [KB83] and Kühnel [Küh95, Thm. 4.13] for more information.)

Exercises

1. Prove the following "generalized Van Kampen–Flores theorem" (Sarkaria [Sar91b]): the simplicial complex

$$T(k_1, \ldots, k_m) \quad := \quad (\sigma^{\le 2k_1})^{\le k_1 - 1} * (\sigma^{\le 2k_2})^{\le k_2 - 1} * \cdots * (\sigma^{\le 2k_m})^{\le k_m - 1}$$

does not embed into $\mathbb{R}^{2k}$, for any partition $k_1+k_2+\cdots+k_m = k+1$.
Similarly (indeed: equivalently!), the complexes $\sigma^r * T(k_1,\ldots,k_m)$ do not embed into $\mathbb{R}^d$ for $d = 2(k_1+k_2+\cdots+k_m-1)+(r+1)$.
(Sarkaria calls these complexes the *Kuratowski complexes*; they include both the Kuratowski minimal nonplanar graphs $K_{3,3}$ and K_5.)

2. Consider a graph G as a 1-dimensional simplicial complex. Prove that G is planar if and only if $\mathrm{ind}_{\mathbb{Z}_2}(G^{*2}_{\Delta}) \leq 2$; that is, Proposition 5.5.1 works perfectly for 1-dimensional simplicial complexes.

3. (Dol'nikov's theorem revisited) Let $\mathcal{P}_0(n) := \{(A,B) : A,B \subseteq [n], A \cap B = \varnothing\}$, and define a partial ordering $\preceq$ on $\mathcal{P}_0$ by inclusion in both components: $(A,B) \preceq (A',B')$ if and only if $A \subseteq A'$ and $B \subseteq B'$. Let $\mathcal{P}(n,\ell) := \{(A,B) \in \mathcal{P}_0(n) : |A \cup B| \geq n-\ell\}$, and let $\mathsf{K}(n,\ell)$ be the order complex of $(\mathcal{P}(n,\ell),\preceq)$. A simplicial $\mathbb{Z}_2$-action on $\mathsf{K}(n,\ell)$ is given by the exchange of the components.

   (a) Check that $\mathsf{K}(n,n-1)$ is isomprphic to $\mathrm{sd}((\sigma^{n-1})^{*2}_{\Delta})$.

   (b) Show that $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}(n,\ell)) \leq \ell$.

   (c) Express $\mathsf{K}(n,n-1)$ as the join of $\mathsf{K}(n,\ell)$ with another suitable simplicial complex, and use Proposition 5.2.2(iii) to verify that $\mathrm{ind}_{\mathbb{Z}_2}(\mathsf{K}(n,\ell)) \geq \ell$, $1 \leq \ell \leq n$.

   (d) Let $\mathcal{S}$ be a set system on $[n]$ and suppose that $c\colon \mathcal{S} \to [m]$ is a proper coloring of the Kneser graph $\mathrm{KG}(\mathcal{S})$. For $A \subseteq [n]$, define $h(A) := \{c(S) : S \in \mathcal{S}, S \subseteq A\}$. Assuming that $\mathrm{cd}_2(\mathcal{S}) > m$, show that the mapping $g\colon (A,B) \mapsto (h(A),h(B))$ is a simplicial $\mathbb{Z}_2$-map of $\mathsf{K}(n,m)$ into $\mathsf{K}(m,m-1)$. Derive Dol'nikov's theorem 3.4.1.

   This proof follows [Mat01b] and the basic idea is due to Kříž [Kri92], [Kri00].

6

Multiple points of coincidence

Up until now, we have been considering spaces with $\mathbb{Z}_2$-actions and theorems saying that under suitable conditions, there exist points $\boldsymbol{x}$ and $\boldsymbol{y}$ with disjoint supports that are mapped to the same point. Here we generalize these considerations on spaces with actions of other groups, most notably the groups $\mathbb{Z}_p$. We will obtain theorems in which images of some p points with disjoint supports are guaranteed to coincide.

6.1 G-spaces

Some spaces posses symmetries other than antipodality: they have groups other than $\mathbb{Z}_2$ acting on them.

For a finite group G, a G-action on a topological space X is a collection $\Phi = (\varphi_g)_{g \in G}$ of homeomorphisms $\varphi_g : X \to X$. The homeomorphism φ_e corresponding to the unit element e of G is the identity id_X, and the composition of these homeomorphisms respects the group operation: $\varphi_g \circ \varphi_h = \varphi_{gh}$ for all $g, h \in G$. (Thus, $g \mapsto \varphi_g$ is a homomorphism of G into the group of homeomorphisms of X; if X is a topological vector space and all the φ_g are linear maps, we have a representation of G in the usual sense.) In the literature, one often writes just gx for $\varphi_g(x)$.

How does our earlier definition of a $\mathbb{Z}_2$-action fit into this general definition? For $G = \mathbb{Z}_2$, the cyclic group $\{0, 1\}$ with addition modulo 2, the homeomorphism assigned to 0 must be the identity, and the homeomorphism assigned to 1 is what was earlier called the $\mathbb{Z}_2$-action ν.

Similarly, let us consider a cyclic group $\mathbb{Z}_n$, which we think of as $\{0, 1, \ldots, n-1\}$ with addition modulo n. A $\mathbb{Z}_n$-action Φ is fully specified by the single homeomorphism φ_1, as $\varphi_i = (\varphi_1)^i$. In this sense, we will mostly write "a $\mathbb{Z}_n$-space (X, ν)," with the action denoted by a lowercase Greek letter, meaning that ν is the homeomorphism corresponding to 1.

We will work exclusively with actions of finite groups, but we can as well state the general definition. For infinite topological groups,[1] we moreover require that $\varphi_g(x)$ depend continuously on both g and x.

[1] A topological group is a group and, at the same time, a Hausdorff topological space, such that the group operation and the inverse are continuous maps $G \times G \to G$ and $G \to G$, respectively.

6.1.1 Definition ($\boldsymbol{G}$-spaces and $\boldsymbol{G}$-maps). *Let G be a topological group and X a topological space. A G-action on X is a collection $\Phi = (\varphi_g)_{g\in G}$ of homeomorphisms $X \to X$, such that $(g,x) \mapsto \varphi_g(x)$ is a continuous map $G \times X \to X$, $\varphi_e = \mathrm{id}_X$, and $\varphi_g \circ \varphi_h = \varphi_{gh}$ for all $g, h \in G$. The pair (X, Φ) is a* **$\boldsymbol{G}$-space**.

If (X, Φ) and (Y, Ψ) are G-spaces, a continuous map $f\colon X \to Y$ is a **$\boldsymbol{G}$-map** *(or* **equivariant map***) if $f \circ \varphi_g = \psi_g \circ f$ for all $g \in G$.*

For $x \in X$, the set $\{\varphi_g(x) : g \in G\}$ is called the *orbit* of x under the G-action Φ, and similarly for a subset $A \subseteq X$. A set $A \subseteq X$ is *invariant* if $\varphi_g(A) = A$ for all $g \in G$.

Free actions. For $\mathbb{Z}_2$-spaces, we have seen the important distinction between free and non-free spaces. We recall that a free $\mathbb{Z}_2$-spaces is one where the single homeomorphism corresponding to 1 has no fixed points. Two ways of generalizing this to actions of larger groups suggest itself: we can require that none of φ_g with $g \neq e$ have a fixed point, or only require that no point be fixed by all φ_g. Both ways lead to interesting notions. We will mostly encounter the former:

6.1.2 Definition. *A G-space (X, Φ) is called* **free** *if no φ_g, $g \neq e$, has a fixed point. Equivalently, for each $x \in X$, the mapping $g \mapsto \varphi_g(x)$ is injective; that is, the orbit of each point is a copy of G.*

The second notion is a *fixed-point free* G-action, where the orbit of each $x \in X$ has at least two points. Our moderate topological means won't allow us to make use of fixed-point free actions that are not free. But in some more advanced applications, they have been employed successfully.

6.1.3 Observation. *Let p be a prime number. Then a $\mathbb{Z}_p$-space (X, ν) is free if and only if ν has no fixed point.*

Indeed, for every k with $1 \leq k < p$, there is some ℓ with $k\ell \equiv 1 \pmod p$, and hence $\nu^k(x) = x$ would imply that $\nu(x) = \nu^{k\ell}(x) = x$.

Examples of group actions. Some of the examples below, especially those with infinite groups, serve just as illustrations, but others (marked by boldface labels) will be important later for combinatorial and geometric applications.

6.1.4 Examples (Group actions).

- **(a)** Let S^1 be the unit circle in the plane and ν the rotation by $\frac{2\pi}{q}$. Then (S^1, ν) is a (free) $\mathbb{Z}_q$-space, for any integer $q > 1$.
- (b) The group $SO(2)$ of all rotations of the plane around the origin also acts on S^1, and we have an example of a (free) $SO(2)$-space.
- (c) More generally, the group $O(n)$ of all isometries of $\mathbb{R}^n$ fixing the origin (corresponding to all orthogonal $n \times n$ matrices with determinant ± 1) acts on S^{n-1} in the obvious way. The action is fixed-point free but not free for $n > 2$. Of course, $O(n)$ acts on $\mathbb{R}^n$ as well, and here the origin is a fixed point.

(d) Since $O(n)$ acts on S^{n-1}, its subgroups $G \subseteq O(n)$ do as well. Such actions are usually called *orthogonal representations* of G, and they have been much studied in the literature. For a slightly exotic example, consider the regular icosahedron

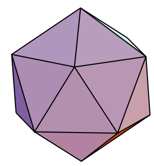

centered at the origin. It is known that the group of symmetries of the icosahedron is A_5 (the noncommutative *alternating group*, consisting of all even permutations of five elements, with composition of permutations as the group operation). Thus, A_5 acts on the icosahedron, and also on its boundary. The latter action is fixed-point free but not free.

(e) In the complex plane, the unit circle S^1 consists of the unit complex numbers: $\{z \in \mathbb{C} : |z| = 1\}$. In this way, S^1 is given a group structure, with complex multiplication as the group operation. Then S^1 is a (free) S^1-space, where the homeomorphism φ_z is given by multiplication by z. Geometrically, multiplication by $z = e^{\mathrm{i}\alpha}$ acts as the rotation of S^1 by the angle α (radians). Thus, this is just a different view of the example (b) with the group of all rotations of the plane around the origin acting on S^1.

(f) *Any* topological group G acts freely on itself by the left multiplication; i.e. $\varphi_g(h) = gh$. The previous example was a special case of this.

(g) New G-spaces can be produced from old ones by *joins*. If (X, Φ) and (Y, Ψ) are G-spaces, then a G-action $\Theta = \Phi * \Psi$ on $X * Y$ is defined by $\theta_g = \varphi_g * \psi_g$. If both Φ and Ψ are free, then the join $\Phi * \Psi$ is free, too. You may want to check that joins of G-maps produce G-maps (Exercise 1). A similar construction can be made for Cartesian products of G-spaces.

(h) The previous abstract example is more clever than it might seem. As we know, the sphere S^3 can be represented as the join $S^1 * S^1$. Taking the rotation by $\frac{2\pi}{q}$ as in (a) on both copies of S^1 and using the join construction in (g), we get a free $\mathbb{Z}_q$-action on S^3. Such an example is by no means obvious. If we consider S^1 as the simplicial complex formed by the perimeter of a regular q-gon, we obtain a triangulated S^3, and the $\mathbb{Z}_q$-action is a simplicial map. Here is an attempt at visualization of the join in $\mathbb{R}^3$. Two hexagons are placed in perpendicular planes, and only the simplex $\{3,4\}*\{1',2'\}$ is shown. Its image under the $\mathbb{Z}_6$-action is $\{4,5\}*\{2',3'\}$ (indicated by dashed lines).

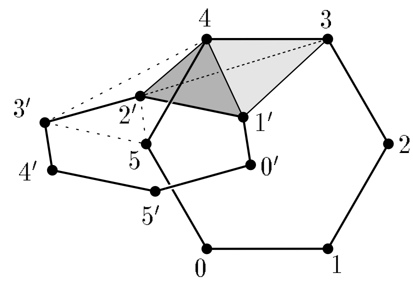

Of course, if we added more simplices to the picture, they would start to intersect; S^3 cannot be embedded in $\mathbb{R}^3$.

(i) In the same way, we get a free $\mathbb{Z}_q$-action on each odd-dimensional sphere S^{2n-1}, using $S^{2n-1} \cong (S^1)^{*n}$. Here is another way of representing the same $\mathbb{Z}_q$-action: regard S^{2n-1} as the unit sphere in $\mathbb{C}^n$, i.e. the set $\{(z_1, \ldots, z_n) \in \mathbb{C}^2 : |z_1|^2 + \cdots + |z_n|^2 = 1\}$, and define the action by $(z_1, \ldots, z_n) \mapsto (\omega z_1, \ldots, \omega z_n)$, where $\omega = e^{2\pi \mathrm{i}/q}$ is a qth root of unity.

(j) It is useful to remember, some negative results, too: the *only* nontrivial group with a *free* action on an even-dimensional sphere S^{2n} is $\mathbb{Z}_2$.[2] Further, it is known that any group G acting freely on some S^n has at most one element of order 2 and every Abelian subgroup of such G is cyclic (equivalently, there is no subgroup $\mathbb{Z}_p \times \mathbb{Z}_p$ with prime p); see e.g. [Hat01, Sec. 1.3] for a part of the proof and references.

(k) For any space X, the symmetric group S_n (all permutations of $[n]$) acts on the nth Cartesian power X^n by permuting the coordinates. Explicitly, for $\pi \in S_n$, the action is $\varphi_\pi(x_1, x_2, \ldots, x_n) = (x_{\pi(1)}, x_{\pi(2)}, \ldots, x_{\pi(n)})$. The subgroups of S_n, such as $\mathbb{Z}_n$, thus act on X^n as well. The same applies to the n-fold joint X^{*n}. These actions are not free but they become free by deleting all fixed points; we will discuss this further when considering deleted joins (and products).

Notes. Actions of groups other than $\mathbb{Z}_2$ on spheres, and the corresponding Borsuk–Ulam type results, appeared soon after Borsuk's paper; Steinlein [Ste85] gives Eilenberg [Eil40] and Hirsch [Hir37], [Hir43] as the earliest such references. They use degree-theoretic considerations or the Lefschetz number. Smith [Smi42], [Smi41], [Smi38] also considered actions of finite groups, but his results mainly concern the structure of the set of fixed points. Many subsequent generalizations asserting the nonexistence of an equivariant map $X \to Y$ relax the conditions on X and Y. For example, since the degree of a map can be defined in (co)homological terms, it suffices that X has the $\mathbb{Z}_2$ (co)homology of S^n and Y that of S^{n-1}. We refer to Steinlein [Ste85] for a detailed bibliography.

[2] This is because the order of a group G having a free action on X must divide the Euler characteristic of X, and the Euler characteristic of S^{2n} is 2.

A basic book on group actions on topological spaces is Bredon [Bre72].
*** mention tom Dieck book???

Exercises

1. If (X_1, Φ_1), (X_2, Φ_2), (Y_1, Ψ_1), and (Y_2, Ψ_2) are G-spaces and $f_1\colon (X_1, \Phi_1) \to (Y_1, \Psi_1)$ and $f_2\colon (X_2, \Phi_2) \to (Y_2, \Psi_2)$ are G-maps, check that $f_1 * f_2\colon (X_1 * X_2, \Phi_1 * \Phi_2) \to (Y_1 * Y_2, \Psi_1 * \Psi_2)$ is a G-map.

6.2 E_nG spaces and the G-index

Much of the theory we have developed for $\mathbb{Z}_2$-spaces, concerning the $\mathbb{Z}_2$-index and the nonexistence of equivariant maps, can be imitated for G-spaces. A large part of this goes through almost without change; we will mainly point out the modifications needed for G-spaces.

As expected, we write $X \xrightarrow{G} Y$ or $X \leq_G Y$ if there is a G-map $X \to Y$. For introducing a G-index, though, we need suitable "yardstick" spaces analogous to the spheres; these are called E_nG spaces.

6.2.1 Definition. *Let G be a finite group and $n \geq 0$. An* **E_nG space** *is a G-space that is*

- *a finite simplicial G-complex (or a finite cell G-complex),*
- *n-dimensional,*
- *$(n-1)$-connected,*
- *and free.*

(Here, similar to the $\mathbb{Z}_2$ case, a **simplicial G-complex** *is a simplicial complex made into a G-space so that all the homeomorphisms φ_g are simplicial maps, and similarly for cell G-complexes.)*

A concrete example of an E_nG space that we will use most often is the $(n+1)$-fold join $G^{*(n+1)}$. As a topological space, this is the join $[m]^{*(n+1)}$, where $m := |G|$ and $[m]$ denotes the m-point discrete space. For example, for $n = 1$, G^{*2} is the complete bipartite graph $K_{m,m}$. Clearly, G^{*n+1} is an n-dimensional simplicial complex. As in Example 6.1.4(f), G acts on itself freely by the left multiplication, and so G^{*n+1} is a free simplicial G-complex. Finally, the $(n-1)$-connectedness follows immediately from Proposition 4.3.5 about the connectivity of joins. With some more work, one can also show by induction that G^{*n+1} is homotopy equivalent to a wedge of a suitable number of n-spheres, from which the $(n-1)$-connectedness follows as well.

We describe other, perhaps simpler, E_nG spaces for the most often considered case $G = \mathbb{Z}_p$. As we know from Example 6.1.4(i), odd-dimensional spheres can be equipped with free $\mathbb{Z}_p$-actions (Example 6.1.4(i)), and so S^{2n-1} with such a $\mathbb{Z}_p$-action can serve as another $E_{2n-1}\mathbb{Z}_p$. For even dimensions, we can take the join of a sphere of dimension one less with one copy of $\mathbb{Z}_p$ (such spaces were used in the first proof of the topological Tverberg theorem by Bárány, Shlosman, and Szücs [BSS81]). We can picture this space as p "tipis" of different heights erected over the sphere S^{2n-1}. As the following picture indicates, this space is homotopy equivalent to a wedge of $(p-1)$ spheres S^{2n} and thus $(2n-1)$-connected.

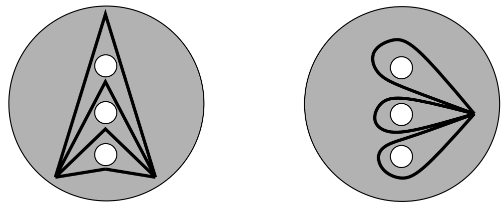

The following lemma shows, among others, that all E_nG spaces are equivalent for our purposes:

6.2.2 Lemma. *Let X be an $(n-1)$-connected G-space and let K be a free finite simplicial G-complex (or a free finite cell G-complex) of dimension at most n. Then $\|\mathsf{K}\| \xrightarrow{G} X$.*

Sketch of proof. The proof is very similar to the proof of Proposition 5.2.2(v): the required G-map is built face-by-face, by induction on the dimension. Having constructed the mapping on the k-skeleton of K, we partition the $(k+1)$-simplices into orbits, we extend the mapping on one simplex in each orbit using k-connectedness, and we transfer this extension to the remaining simplices via the G-action. Here we need that the simplices in each orbit have disjoint relative interiors, but if the relative interior of $\varphi_g(\operatorname{int}\sigma)$ intersected the relative interior of σ, then we would have $\varphi_g(\sigma) = \sigma$ (as φ_g is simplicial and bijective) and σ would contain a point fixed by φ_g.

6.2.3 Definition (G-index). *For a G-space X, we define*

$$\operatorname{ind}_G(X) := \min\{n : X \xrightarrow{G} E_nG\}.$$

(Here E_nG can be any E_nG space, since any of them G-maps into any other.)

The properties of the $\mathbb{Z}_2$-index listed in Proposition 5.2.2 generalize without change. For convenience, we list them again here; we also add Sarkaria's inequality.

6.2.4 Proposition (Properties of the G-index). *Let G be a nontrivial finite group ($|G| > 1$).*

(i) $\operatorname{ind}_G(X) > \operatorname{ind}_G(Y)$ *implies* $X \stackrel{G}{\not\longrightarrow} Y$.

(ii) $\operatorname{ind}_G(E_nG) = n$ *(for any E_nG space).*

(iii) $\operatorname{ind}_G(X * Y) \leq \operatorname{ind}_G(X) + \operatorname{ind}_G(Y) + 1$.

(iv) *If X is $(n-1)$-connected, then* $\operatorname{ind}_G(X) \geq n$.

(v) *If X is a* **free** *simplicial G-complex (or free cell G-complex) of dimension n, then* $\operatorname{ind}_G(X) \leq n$.[3]

(vi) *(Sarkaria's inequality) If L is a finite simplicial G-complex and L_0 is an invariant subcomplex of it, then* $\operatorname{ind}_G(\mathsf{L}_0) \geq \operatorname{ind}_G(\mathsf{L}) - \operatorname{ind}_G(\Delta(\mathsf{L} \setminus \mathsf{L}_0)) - 1$.

Part (i) is obvious, (iii) follows from the fact that G^{*n+1} is an E_nG space, (iv) and (v) are consequences of Lemma 6.2.2 (of course, (iv) also needs (ii)), and (vi) is proved exactly like Theorem 5.6.2. The hardest part is the innocent-looking (ii), which requires a new theorem of a Borsuk–Ulam type.

6.2.5 Theorem (A "Borsuk–Ulam" theorem for G-spaces). *There is no G-map of an E_nG space into an $E_{n-1}G$ space.*

We postpone the proof a little, and we comment on the role of the groups $\mathbb{Z}_p$. First, we observe that if H is a subgroup of G, then any G-space can also be regarded as an H-space (and a G-map as an H-map). By inspecting the above proposition, we see that it never makes any reference to the properties of G (except for the nontriviality), and so, if we use only these tools for bounding the index, we lose nothing by restricting ourselves to a nontrivial subgroup. In fact, sometimes we might gain, since it can happen that a G-action is not free but the action of some subgroup H is free. It is not hard to show that every (nontrivial) finite group contains a subgroup isomorphic to $\mathbb{Z}_p$ for a prime p. Therefore, when considering free actions, it is usually sufficient to consider only $\mathbb{Z}_p$-actions. This happens, for instance, in the following proof.

Sketch of proof of Theorem 6.2.5. (Specialized to $G = \mathbb{Z}_2$, this is also another proof of the Borsuk–Ulam theorem.) Exceptionally, in this proof we have to assume familiarity with the basics of simplicial homology.

As was just noted above, it is sufficient to consider the case $G = \mathbb{Z}_p$ with prime p.

For concreteness, let us work with the $E_n\mathbb{Z}_p$ space $(\mathbb{Z}_p)^{*n+1}$. Let $\mathsf{K} := (\mathbb{Z}_p)^{*n+1}$ and let $\mathsf{L} := (\mathbb{Z}_p)^{*n}$. This L can be identified with a subcomplex of K (corresponding to the first n factors in the $(n+1)$-fold join); let $i\colon \mathsf{L} \to \mathsf{K}$ be the inclusion map.

For contradiction, we suppose that there is a $\mathbb{Z}_p$-map $f\colon \mathsf{K} \to \mathsf{L}$. First we need to make f into a simplicial map; more precisely, we need that there is a sufficiently fine subdivision $\tilde{\mathsf{K}}$ of K and a simplicial $\mathbb{Z}_p$-map $\tilde{f}\colon \tilde{\mathsf{K}} \to \mathsf{L}$. This is done using a standard procedure (simplicial approximation theorem; see e.g. [Hat01, Theorem 2C.1]); one has to be a little careful so that the simplicial approximation remains a $\mathbb{Z}_p$-map, but this is not a problem.

The composed map $g := i \circ \tilde{f}\colon \tilde{\mathsf{K}} \to \mathsf{K}$ is a simplicial $\mathbb{Z}_p$-map. We analyze its *Lefschetz number* in two ways and reach a contradiction.

First we consider the level of chain groups. The simplicial map $g\colon \tilde{\mathsf{K}} \to \mathsf{K}$ induces maps $g_{\#k}\colon C_k(\tilde{\mathsf{K}}) \to C_k(\tilde{\mathsf{K}})$, where $C_k(\tilde{\mathsf{K}}) = C_k(\tilde{\mathsf{K}}, \mathbb{Q})$ is the k-dimensional

[3] As in Proposition 5.2.2, this holds for all paracompact Y of dimension at most n.

chain group with rational coefficients (g goes into K but every k-simplex in K is written as the sum of the k-simplices in $\tilde{\mathsf{K}}$ subdividing it). The Lefschetz number on the level of chain maps is

$$\Lambda(g) = \sum_{k \geq 0} (-1)^k \operatorname{trace}(g_{\#k}).$$

Since we are working with rational coefficients, the $C_k(\tilde{\mathsf{K}})$ are vector spaces and the $g_{\#k}$ are linear endomorphisms, and so trace is the trace of a linear map in the usual sense.

We consider the usual basis of $C_k(\tilde{\mathsf{K}})$ made of all chains e_σ, where $\sigma \in \tilde{\mathsf{K}}$ has dimension k and e_σ is 1 on σ and 0 elsewhere. Expressing $\operatorname{trace}(g_{\#k})$ with respect to this basis, we see that since g is a $\mathbb{Z}_p$-map, σ gives the same contribution as the other $p-1$ simplices in its orbit (here we use that the simplices in each orbit are all distinct). Therefore, $\operatorname{trace}(g_{\#k})$ is divisible by p, and so is $\Lambda(g)$.

Now we consider $\Lambda(g)$ on the level of homology groups. The map g induces maps $g_{*k}\colon H_k(\mathsf{K}, \mathbb{Q}) \to H_k(\mathsf{K}, \mathbb{Q})$ in homology, and by the *Hopf trace formula*, the Lefschetz number equals

$$\Lambda(g) = \sum_{k \geq 0} (-1)^k \operatorname{trace}(g_{*k}).$$

Since K is $(n-1)$-connected, we have $H_k(\mathsf{K}, \mathbb{Q}) = 0$ for $1 \leq k \leq n-1$, and so the only contribution to $\Lambda(g)$ may come from dimensions 0 and n. But g_{*n} is trivial, since it is the composition $i_{*n} \circ \tilde{f}_{*n}$, and so it goes through the homology group $H_n(\mathsf{L}, \mathbb{Q})$ which is 0 because L is $(n-1)$-dimensional. It follows that $\Lambda(g) = 1$, which contradicts the previous calculation and shows that the $\mathbb{Z}_p$-map $f\colon \|K\| \to \|L\|$ is impossible. (From the first part of the proof, we can actually learn something about actual (existing) $\mathbb{Z}_p$-maps of a (triangulable) $\mathbb{Z}_p$-space into itself: any such map has Lefschetz number divisible by p.)

The following consequence of Proposition 6.2.4, which does not mention the G-index, has been often quoted and used in the literature:

6.2.6 Theorem (Dold's theorem [Dol83]). *Let X be an n-connected G-space and let Y be a free G-space of dimension at most n (it may be a simplicial G-complex, a cell G-complex, or even an arbitrary paracompact space). Then $X \overset{G}{\not\to} Y$.*

Notes. Krasnosel'skiĭ's notion of genus (equivalent to our $\mathbb{Z}_2$-index) was extended to actions of more general groups by Švarc [Šva57], [Šva62].

There are more advanced results, whose proofs or even reasonably general formulation are beyond our scope, which can establish $X \overset{G}{\not\to} Y$ with the G-action on Y being fixed-point free but not necessarily free. One useful result, which can be formulated easily, is the following theorem of Volovikov [Vol96]:

Let $G := \mathbb{Z}_p \times \mathbb{Z}_p \times \cdots \times \mathbb{Z}_p$ be the product of finitely many copies of $\mathbb{Z}_p$, with p prime. Let X and Y be fixed-point free G-spaces such that $\tilde{H}^i(X, \mathbb{Z}_p) = 0$ for all $i \leq n$ (reduced cohomology groups with $\mathbb{Z}_p$-coefficients) and Y is finite-dimensional and an n-dimensional cohomology sphere over $\mathbb{Z}_p$ (the cohomology ring with $\mathbb{Z}_p$ coefficients is isomorphic to that of S^n). Then $X \stackrel{G}{\not\to} Y$.

In particular, there is no G-map of an n-connected X into S^n, provided that the actions are fixed-point free. Similar results have been obtained, in varying degrees of generality, by Özaydin [Öza] (in an unpublished manuscript) and later independently by Sarkaria [Sar00]. A detailed completion and exposition of Sarkaria's argument was given by de Longueville [dL99]. The proofs rely on advanced topological methods (cohomology and characteristic classes of vector bundles).

Cohomological ideal-valued index. Production of results similar to the just mentioned theorem can be "mechanized" using a cohomological ideal-valued index of Fadell & Husseini [FH88]. This index theory appears very useful for combinatorial and geometric applications. It can be seen as a generalization of the idea of the cohomological proof of the Borsuk–Ulam theorem mentioned in the notes to Section 2.1. The index $\mathrm{Ind}_G(X)$ of a G-space (X, Φ) is not a single number, but rather an ideal in a ring R_G. This ring is the cohomology ring of a certain space constructed from G and, for finite G, it can usually be represented as a polynomial ring. A G-map $(X, \Phi) \to (Y, \Psi)$ implies the containment $\mathrm{Ind}_G(X) \subseteq \mathrm{Ind}_G(Y)$, and so the existence of a G-map can be excluded whenever this inclusion doesn't hold. This index is finer than the numerical G-index considered in this chapter, and it also gives results for fixed-point free actions. On the other hand, its computation requires the knowledge of certain cohomology rings and their maps, which may not be easy to obtain. A short introduction to this theory with several impressive applications and a few ready-made recipes for computing $\mathrm{Ind}_G(X)$ in some common cases was provided by Živaljević [Živ98].

Finally, the *equivariant obstruction theory* is another powerful tool (again requiring more advanced knowledge of algebraic topology) for attacking the question whether $X \xrightarrow{G} Y$ or not. Sometimes it yields the nonexistence of a G-map and sometimes, unlike the index theories, it allows one to prove the *existence* of a G-map $X \to Y$ (without explicitly constructing it). For an application, the existence of a G-map is usually disappointing but at least it identifies a dead end. Equivariant obstruction theory deals with the following question: Given an equivariant map f defined on the n-skeleton of a simplicial G-complex (or cell G-complex), is there an equivariant map defined on the $(n+1)$-skeleton that agrees with f on the $(n-1)$-skeleton? (In other words, we want to extend f from the $(n-1)$-skeleton to the $(n+1)$-skeleton, knowing that extension to the n-skeleton is possible.) The answer is yes if and only if certain cohomology class (the "obstruction") is zero. Since there can be many choices for the extension in each step, the method doesn't seem to provide a generally efficient algorithm for deciding whether $X \xrightarrow{G} Y$, even if we can evaluate the required cohomology classes. In many concrete cases it works nicely, though. For a first impression of the method, one can consult [Živ98], which also provides references for a deeper study.

Exercises

1. Prove by induction on n that $[m]^{*n}$ is homotopy equivalent to a wedge of $(n-1)$-dimensional spheres. How many spheres are there?

6.3 Deleted joins and deleted products

In the subsequent applications, which are mostly generalizations of problems we have encountered earlier, we construct G-spaces X and Y and then use the G-index for showing $X \stackrel{G}{\not\to} Y$. Here X and Y are usually suitable p-fold deleted joins or deleted products, and in this section we discuss these constructions. Unlike for twofold joins and products, for p-fold ones there are various possibilities as to which points should be deleted. For example, from the product X^3, we can delete all points (x,x,x), where all the three components coincide, or alternatively the points where at least two coordinates coincide. What needs to be deleted is usually dictated by the application. Here is the general definition, of which we will actually use only a few special cases.

6.3.1 Definition. *Let $n \geq k \geq 2$ be given integers (we will mostly encounter the cases $k=n$ and $k=2$). Call an n-tuple $(x_1, x_2, \ldots, x_n)$* ***k-wise distinct*** *if no k among the x_i are equal.*

The ***n-fold k-wise deleted product*** *of a space X is*

$$X^n_{\Delta(k)} := X^n \setminus \{(x_1, x_2, \ldots, x_n) \in X^n : (x_1, \ldots, x_n) \text{ not } k\text{-wise distinct}\}.$$

The ***n-fold k-wise deleted join*** *of X is*

$$X^{*n}_{\Delta(k)} := X^{*n} \setminus \{\tfrac{1}{n}x_1 + \tfrac{1}{n}x_2 + \cdots + \tfrac{1}{n}x_n : (x_1, \ldots, x_n) \text{ not } k\text{-wise distinct}\}.$$

For a simplicial complex K, the ***n-fold k-wise deleted join*** *of K is*

$$\mathsf{K}^{*n}_{\Delta(k)} := \{F_1 * F_2 * \cdots * F_n \in \mathsf{K}^{*n} : (F_1, F_2, \ldots, F_n) \text{ } k\text{-wise disjoint}\},$$

where an n-tuple $(F_1, F_2, \ldots, F_n)$ of sets is ***k-wise disjoint*** *if every k among the F_i have empty intersection.*

For $k=n$, we write only X^n_Δ for $X^n_{\Delta(n)}$, X^{*n}_Δ for $X^{*n}_{\Delta(n)}$, and K^{*n}_Δ for $\mathsf{K}^{*n}_{\Delta(n)}$.

So the 2-wise deleted joins and products are the "most deleted" (smallest) while the n-wise deleted ones are the "least deleted" (largest).

On all these deleted joins and products, the symmetric group S_n acts by permuting the coordinates. We will consider the action of the cyclic subgroup

$\mathbb{Z}_n$ generated by the cyclic shift to the left, namely by the permutation ν: $1 \mapsto 2$, $2 \mapsto 3$, ..., $n{-}1 \mapsto n$, $n \mapsto 1$. Explicitly, on the deleted product, ν acts by

$$\nu\colon\ (x_1, x_2, \ldots, x_n)\ \longmapsto\ (x_2, x_3, \ldots, x_n, x_1),$$

and on the deleted join, it acts by

$$\nu\colon\ t_1x_1 + t_2x_2 + \cdots + t_nx_n\ \longmapsto\ t_2x_2 + t_3x_3 + \cdots + t_nx_n + t_1x_1.$$

Free actions. For 2-wise deleted joins and products, where no two coordinates of points coincide, the whole S_n-action is free.

On the other hand, for n-wise deleted n-fold products and joins, the S_n-action is not free and the $\mathbb{Z}_n$-action ν is free *if (and only if) $n = p$ is a prime.* Indeed, if p is a prime, then by Observation 6.1.3, it suffices to verify that ν has no fixed point, and this is obvious since if $(x_2, x_3, \ldots, x_n, x_1) = (x_1, x_2, \ldots, x_n)$, then $x_1 = x_2 = \cdots = x_n$. Moreover, as is not difficult to check, this is the only case (up to a renumbering of the coordinates) when a nontrivial subgroup of S_n acts freely on an n-wise deleted n-fold product or join of a space or simplicial complex with at least two points (Exercise 1).

We need deleted joins and products of spaces only for the case $X = \mathbb{R}^d$ and $k = n$. Now we calculate the $\mathbb{Z}_p$-indices in that case.

6.3.2 Proposition (Deleted products and deleted joins of $\mathbb{R}^d$). *Let p be a prime and let $d \geq 1$. Then*

$$\mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^p_\Delta) \leq d(p{-}1){-}1$$

and

$$\mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^{*p}_\Delta) \leq (d{+}1)(p{-}1){-}1.$$

Proof. We construct a $\mathbb{Z}_p$-map $g\colon (\mathbb{R}^d)^p_\Delta \to S^{d(p-1)-1}$, where $S^{d(p-1)-1}$ is equipped with a suitable free $\mathbb{Z}_p$-action.

Let us interpret $\mathbb{R}^{d\times p} = \left(\mathbb{R}^d\right)^p$ as the space of matrices $(x_{ij})_{i=1}^{d}{}_{j=1}^{p}$ with d rows and p columns. The $\mathbb{Z}_p$-action is the cyclic shift of the columns. The elements of $(\mathbb{R}^d)^p_\Delta$ are all matrices of this form except for those with all columns being equal. For instance, for $d = 1$ and $p = 3$ we get the 3-dimensional Euclidean space with the diagonal line $\{x_1 = x_2 = x_3\}$ removed. First we consider the orthogonal projection g_1 of $\mathbb{R}^{d\times p}$ on the $d(p{-}1)$-dimensional subspace L perpendicular to the diagonal. In coordinates, L is the subspace consisting of all $d{\times}p$ matrices with zero row sums, and g_1 maps a matrix $X = (x_{ij})$ to the matrix

$$g_1(X) = \left(x_{ij} - \frac{1}{p}\sum_{k=1}^{p} x_{ik}\right)_{ij};$$

that is, the average of all columns is subtracted from each column. We see that $g_1(X)$ is the zero matrix O if and only if each column of X equals the average of all columns; i.e. if all columns of X are equal. Therefore, g_1 provides a (surjective) $\mathbb{Z}_p$-map $(\mathbb{R}^d)^p_\Delta \to L\setminus\{O\}$. For instance, for $d = 1$ and $p = 3$, the map g_1 is the orthogonal projection onto the plane $x_1 + x_2 + x_3 = 0$.

We set $g(X) := \frac{g_1(X)}{\|g_1(X)\|}$. The range of g is the unit sphere $S(L)$ in L, which can be identified with $S^{d(p-1)-1}$. Here is a geometric illustration for $p = 3$ and $d = 1$:

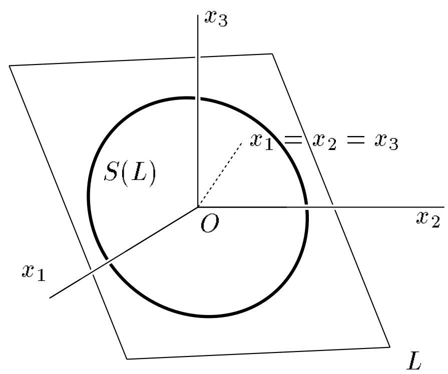

Clearly, g is a $\mathbb{Z}_p$-map, and we have proved the first part of the proposition. As for the deleted join, we construct a $\mathbb{Z}_p$-map $h\colon (\mathbb{R}^d)^{*p}_\Delta \to (\mathbb{R}^{d+1})^p_\Delta$, generalizing the proof of Lemma 5.4.5 in a straightforward manner. As in that proof, we consider the deleted join of a bounded set, say B^d, instead of $\mathbb{R}^d$. Then we place the copies of B^d into $(\mathbb{R}^{(d+1)})^p$ using the embeddings $\psi_1, \ldots, \psi_p$, where $\psi_i(\boldsymbol{x})$ has $(1, x_1, x_2, \ldots, x_d)$ in the ith block of coordinates and 0s elsewhere. The mapping $h\colon (B^n)^{*p}_\Delta \to (\mathbb{R}^{d+1})^p_\Delta$ is given by

$$t_1\boldsymbol{x}_1 + t_2\boldsymbol{x}_2 + \cdots + t_p\boldsymbol{x}_p \longmapsto t_1\psi_1(\boldsymbol{x}_1) + t_2\psi_2(\boldsymbol{x}_2) + \cdots + t_p\psi_p(\boldsymbol{x}_p).$$

It is clearly a $\mathbb{Z}_p$-map, it goes into the deleted product as it should, and continuity follows by a slight generalization of the considerations in the proof of Proposition 4.2.4 about a geometric representation of joins.

With a little more work, it can be shown that $(\mathbb{R}^d)^p_\Delta \simeq S^{d(p-1)-1}$, and so $\mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^p_\Delta)$ actually equals $d(p-1)-1$ (Exercise 2). Similarly, $\mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^{*p}_\Delta) = (d+1)(p-1)-1$.

Warning. For general n and k, the topology of the deleted product $(\mathbb{R}^d)^n_{\Delta(k)}$ can be quite complicated. Based on some special cases and on an analogy with the deleted join of a simplex, $(\sigma^d)^{*n}_{\Delta(k)}$ (which is homotopy equivalent to a wedge of $((d+1)(k-1)-1)$-spheres; see Exercise 6.7.1), one might be tempted to believe that $(\mathbb{R}^d)^n_{\Delta(k)}$ is homotopy equivalent to a wedge of $(d(k-1)-1)$-spheres (as is asserted [Sar91a]). The truth is much subtler, though: while it can be shown that $(\mathbb{R}^d)^n_{\Delta(k)}$ is $(d(k-1)-2)$-connected, it can also have nonzero homology in various higher dimensions.

We conclude this section with generalization of version (1.1) of the Borsuk–Ulam theorem; we recall that (1.1) asserts the existence of an $\boldsymbol{x}$ with $f(\boldsymbol{x}) = f(-\boldsymbol{x})$ for any continuous $f\colon S^n \to \mathbb{R}^n$.

6.3.3 Theorem (On p-fold coincidence points). *Let (X, ν) be a $\mathbb{Z}_p$-space with $\mathrm{ind}_{\mathbb{Z}_p}(X) \geq d(p-1)$, where p is a prime. Then for any continuous map $f: X \to \mathbb{R}^d$ there exists $x \in X$ such that $f(x) = f(\nu(x)) = f(\nu^2(x)) = \cdots = f(\nu^{p-1}(x))$.*

Proof. Suppose that there is no such $x \in X$. Then the map

$$x \longmapsto \Big(f(x), f(\nu(x)), \ldots, f(\nu^{p-1}(x))\Big)$$

is a $\mathbb{Z}_p$-map of X into the deleted product $(\mathbb{R}^d)^p_\Delta$, which yields $\mathrm{ind}_{\mathbb{Z}_p}(X) \leq \mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^p_\Delta) \leq d(p-1)-1$.

Notes. The space $X^n_{\Delta(2)}$ is sometimes called the nth *(ordered) configuration space* of X, since it models configurations of n distinct (and distinguishable) particles in X, and it is a classical object of study. For $X = \mathbb{C} \cong \mathbb{R}^2$, $\mathbb{C}^n_{\Delta(2)}$ is known as the *pure braid space*. Lot of work has been devoted to the topological properties of the complement of the zero set of various systems of polynomials; see Vassiliev [Vas92] for interesting and advanced results.

*** mention Fadell-Husseni book

The topology of the deleted products $(\mathbb{R}^d)^n_{\Delta(k)}$ for $d = 1$ and $d = 2$ has been investigated by Björner and Welker [BW95] (for $d = 1$, $(\mathbb{R})^n_{\Delta(k)}$ is known as the *k-equal manifold*). Their method generalizes easily to arbitrary d and allows one to describe the cohomology in concrete cases, although obtaining general formulas seems very complicated.

Exercises

1. Let X be a topological space with at least two points.

   (a) Show that if n is not a prime then the $\mathbb{Z}_n$-action on X^n_Δ generated by the cyclic shift by one position left is not free.

   (b) More generally, show that if G is a nontrivial subgroup of S_n whose action on X^n_Δ is free then $n = p$ is a prime and G is a cyclic group isomorphic to $\mathbb{Z}_p$ generated by a cyclic shift, after a suitable renumbering of the coordinates.

2. Show that $(\mathbb{R}^d)^p_\Delta$ and $S^{d(p-1)-1}$ are homotopy equivalent. (Use the map g in the proof of Proposition 6.3.2.)

3. For $p = 3$ and $d = 1$, the sphere $S(L)$ in the proof of Proposition 6.3.2 is isometric to S^1. Is it true that the cyclic shift action ν on $S(L)$ inherited from $\mathbb{R}^3$ is equal to the rotation of $S(L)$ by $\frac{2\pi}{3}$?

4. (A Lusternik–Schnirelmann-type theorem for $\mathbb{Z}_p$-actions) Let (X, ν) be a $\mathbb{Z}_p$-space (assume that X is a metric space if it helps) with $\mathrm{ind}_{\mathbb{Z}_p}(X) \geq d(p-1)$, where p is a prime, and let $A_1, A_2, \ldots, A_{d+1}$ be closed sets covering X. Show that there is an index i and a point $x \in X$ such that $\{x, \nu(x), \ldots, \nu^{p-1}(x)\} \subseteq A_i$.

6.4 Necklace for many thieves

We consider the necklace problem from Section 3.2 but with q thieves. We only deal with the continuous version here (the discrete version is proved from the continuous version by a simple combinatorial argument). The following theorem formally states that $q(d-1)$ cuts suffice for q thieves.

6.4.1 Theorem (Continuous necklace with many thieves; Alon [Alo87]). *Let $\mu_1, \mu_2, \dots, \mu_d$ be continuous probability measures on $[0,1]$, let $q \geq 2$, and set $N = d(q-1)$. Then there exists a partition of $[0,1]$ into $N+1$ intervals $I_1, I_2, \dots, I_{N+1}$ by N cuts and a partition of the index set $[N+1]$ into subsets $T_1, T_2, \dots, T_q$ such that*

$$\sum_{j \in T_k} \mu_i(I_j) = \tfrac{1}{q} \quad \text{for } i = 1, 2, \dots, d \text{ and } q = 1, 2, \dots, q.$$

Proof. In the subsequent topological argument, we will need to assume that the number of thieves q is a prime. Unlike the topological Tverberg theorem, say, the non-prime cases follow from the result for all prime q by a simple direct argument; see Exercise 1.

From now on, q is a prime. Consider an arbitrary division of $[0,1]$ among q thieves: let $I_1, I_2, \dots, I_{N+1}$ be a partition of the interval $[0,1]$ into $N+1$ intervals (numbered from left to right), and let $T_1, T_2, \dots, T_q$ be a partition of $[N+1]$. We encode such division by a point of the deleted join $\|(\sigma^N)^{*q}_{\Delta(2)}\|$; this is the key step.

Let us regard σ^N as the "standard simplex" in $\mathbb{R}^{N+1}$:

$$\sigma^N = \{\boldsymbol{x} \in \mathbb{R}^{N+1} : x_j \geq 0,\ x_1 + x_2 + \cdots + x_{N+1} = 1\}.$$

Each of the $N+1$ vertices of σ^N lies on one of the coordinate axes, and so the vertex set can be identified with $[N+1]$.

A point of the deleted join $\|(\sigma^N)^{*q}_{\Delta(2)}\|$ has the form $t_1\boldsymbol{x}_1 + t_2\boldsymbol{x}_2 + \cdots + t_q\boldsymbol{x}_q$. First we determine the coefficients t_k from the given division: t_k is the total length of intervals assigned to the kth thief, i.e.

$$t_k := \sum_{j \in T_k} \mathrm{length}(I_j).$$

Next, we define $\boldsymbol{x}_k$. If $t_k = 0$, then $\boldsymbol{x}_k$ does not matter in the join, so assume $t_k > 0$. We set

$$(\boldsymbol{x}_k)_j := \begin{cases} \frac{1}{t_k}\mathrm{length}(I_j) & \text{for } j \in T_k \\ 0 & \text{for } j \notin T_k. \end{cases}$$

In other words, we consider the intervals going to the kth thief and we blow them up, all in the same ratio, so that they fill up the whole interval $[0,1]$, while the other intervals shrink to zero length. Here is an example for $N = 6$, $q = 3$, and $i = 2$:

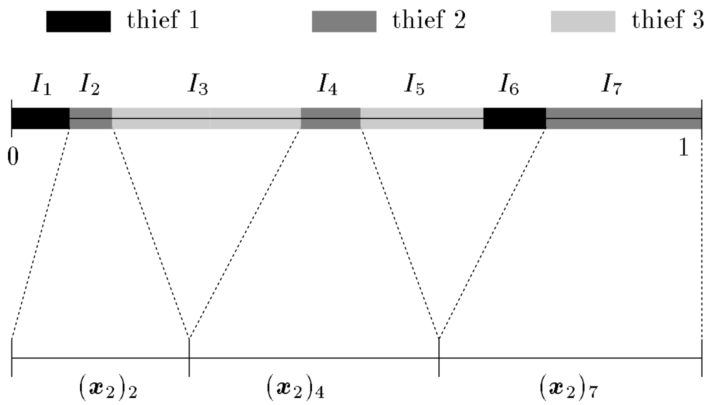

Note that $V(\mathrm{supp}(\boldsymbol{x}_k)) \subseteq T_k$, and so the $\boldsymbol{x}_k$ have pairwise disjoint supports.

Conversely, given any point $\boldsymbol{z} = t_1\boldsymbol{x}_1 + \cdots + t_q\boldsymbol{x}_q \in \|(\sigma^N)^{*q}_{\Delta(2)}\|$, we can determine the lengths of the intervals $I_1, \ldots, I_{N+1}$ uniquely, and we can also find the assignments of the intervals of nonzero lengths to the thieves: T_k consists of the indices of the vertices of $\mathrm{supp}(\boldsymbol{x}_k)$. The assignment of the intervals of zero length is not unique. But what *is* unique is the function $f\colon \|(\sigma^N)^{*q}_{\Delta(2)}\| \to (\mathbb{R}^d)^q$ expressing the gains of the thieves. Namely, we put

$$f(\boldsymbol{z})_{i,k} := \sum_{j \in T_k} \mu_i(I_j).$$

It can be verified that f is continuous, and obviously it is a $\mathbb{Z}_q$-map. If there were no division as claimed in the theorem, f would miss the diagonal in $(\mathbb{R}^d)^q$, and so we would get an equivariant map

$$f\colon\ \|(\sigma^N)^{*q}_{\Delta(2)}\| \ \longrightarrow\ (\mathbb{R}^d)^q_\Delta.$$

This is impossible, since the $\mathbb{Z}_q$-index of the left-hand side is N while that of the right-hand side is $(d+1)(q-1)-1 = N-1$.

Notes. Alon's proof [Alo87] of the necklace theorem for many thieves uses a different encoding of the divisions and relies on the Borsuk–Ulam type result of Bárány, Shlosman, and Szücs [BSS81] mentioned in the notes to Section 6.5. The presented proof basically follows Vućić and Živaljević [VŽ93] (they assume, w.l.o.g., that one of the μ_i is the Lebesgue measure on $[0,1]$, and they construct a $\mathbb{Z}_q$-map into the deleted join $(\mathbb{R}^{d-1})^{*q}_\Delta$ instead of deleted product). They also give a lower bound for the number of fair divisions for "generic" necklaces (where no fair splitting is possible with fewer than $d(q-1)$ cuts), by the method shown in Section 6.6 below for Tverberg partitions.

Exercises

1. Suppose that the statement of Theorem 6.4.1 holds with $q = q_1$ and also with $q = q_2$. Show that it holds for $q = q_1q_2$, too.

6.5 The topological Tverberg theorem

Radon's theorem 5.3.1 states that any $d+2$ points in $\mathbb{R}^d$ can be divided into two parts with intersecting convex hulls. Tverberg's theorem is a generalization of this statement, where we want not only two disjoint subsets with intersecting convex hulls but r of them.

It is not too difficult to show that for every d and r, there exists a $T = T(d, r)$ such that for any set A of T points in $\mathbb{R}^d$ can be divided into r pairwise disjoint subsets $A_1, A_2, \ldots, A_r$ with $\bigcap_{i=1}^r \operatorname{conv}(A_i) \neq \emptyset$ (Exercise 1). It is much harder to establish the tight bound for $T(r, d)$, as stated in the next theorem.

6.5.1 Theorem (Tverberg's theorem [Tve66]). *For any $d \geq 1$ and $r \geq 2$, any set of $(d+1)(r-1)+1$ points in $\mathbb{R}^d$ can be partitioned into r pairwise disjoint subsets $A_1, \ldots, A_r$ in such a way that $\operatorname{conv}(A_1) \cap \cdots \cap \operatorname{conv}(A_r) \neq \varnothing$ (call such a partition a* **Tverberg partition***).*

Let us examine some special cases first. As was remarked above, the case $r = 2$ is Radon's theorem. For $d = 1$, we have $2r-1$ points on the real line, say $x_1 \leq x_2 \leq \cdots \leq x_{2r-1}$. Then we can choose $A_i := \{x_i, x_{2r-i}\}$ for $1 \leq i \leq r-1$, and $A_r = \{x_r\}$. In fact, if the points x_i are all distinct, then this is the only suitable partition! Here is an example for $d = 2$ and $r = 3$, showing two possible Tverberg partitions of a 7-point set (can you find other partitions?):

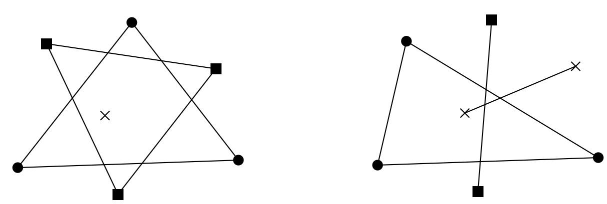

The reader is invited to check that of $N+1$ points with $N = (d+1)(r-1)$ is the best possible (for all r and d); see Exercise 2.

We will not prove Tverberg's theorem here—instead, we prove a topological version which implies Tverberg's theorem in the case where r is a prime.

6.5.2 Theorem (Topological Tverberg theorem; Bárány, Shlosman & Szücs [BSS81]). *Let p be a prime, let $d \geq 1$ be arbitrary, and put $N = (d+1)(p-1)$. For every continuous map*

$$f\colon \|\sigma^N\| \longrightarrow \mathbb{R}^d$$

there exist p pairwise disjoint faces $F_1, \ldots, F_p \subseteq \sigma^N$ whose images under f intersect:

$$f(\|F_1\|) \cap f(\|F_2\|) \cap \cdots \cap f(\|F_p\|) \neq \varnothing.$$

It seems likely that this theorem remains true for all p, not only primes, but so far nobody has managed to prove this. It has been verified for all prime powers, though.

Proof. This is very similar to the proof of the topological Radon theorem; the only difference is that we work with p-fold joins.

Suppose that there is an f violating the theorem; that is, there are no pairwise disjoint faces $F_1, F_2, \ldots, F_p$ with all $f(\|F_i\|)$ intersecting. We consider the p-fold join f^{*p}, and we regard it as a map from the p-fold 2-wise deleted join:

$$f^{*p}\colon \|(\sigma^N)^{*p}_{\Delta(2)}\| \longrightarrow (\mathbb{R}^d)^{*p}_{\Delta}.$$

The fact that this map indeed goes into the deleted join exactly translates the condition on f above.

Note how the problem itself determines what kind of deleted joins we should use: we deal with *pairwise disjoint* faces, and so we use the 2-wise deleted join on the left-hand side. We assume that no p images coincide, and so the join on the right-hand side is p-wise deleted (only the points with all components equal are removed).

Automatically, f^{*p} is a continuous $\mathbb{Z}_p$-map. We know that $\mathrm{ind}_{\mathbb{Z}_p}((\mathbb{R}^d)^{*p}_{\Delta}) \leq (d{+}1)(p{-}1) - 1$ (Proposition 6.3.2), and so it remains to calculate the $\mathbb{Z}_p$-index of the left-hand side. This is again similar to the case $p = 2$ handled in connection with the topological Radon theorem. Analogous to Lemma 5.4.2, we have

6.5.3 Lemma. *Let* K *and* L *be simplicial complexes. Then*

$$(\mathsf{K} * \mathsf{L})^{*p}_{\Delta(2)} = \mathsf{K}^{*p}_{\Delta(2)} * \mathsf{L}^{*p}_{\Delta(2)}.$$

Proof. Clear!

6.5.4 Corollary. *We have* $\mathrm{ind}_{\mathbb{Z}_p}((\sigma^n)^{*p}_{\Delta(2)}) = n$.

Proof. This time we have

$$(\sigma^n)^{*p}_{\Delta(2)} = ((\sigma^0)^{*(n+1)})^{*p}_{\Delta(2)} = ((\sigma^0)^{*p}_{\Delta(2)})^{*(n+1)} = [p]^{*(n+1)}.$$

In Section 6.2, we saw that $[p]^{*(n+1)}$ is $(n{-}1)$-connected—in fact, it is an $E_n\mathbb{Z}_p$ space (if we identify $[p]$ with $\mathbb{Z}_p$).

This also concludes the proof of the topological Tverberg theorem.

The space $[p]^{*(n+1)}$ is quite important; we used it as an $E_n\mathbb{Z}_p$ space, here it turned up as the deleted join of a simplex, and we will meet it several more times. From a combinatorial point of view, the maximal simplices can be regarded as the edges of the complete $(n{+}1)$-partite hypergraph on $n{+}1$ classes of size p each. In the picture, $n = 2$, $p = 4$, and only 3 edges are drawn as a sample:

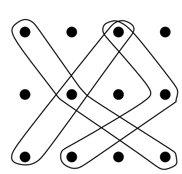

The isomorphism of the complex $[p]^{*(n+1)}$ with the deleted join $(\sigma^n)^{*p}_{\Delta(2)}$ is quite intuitive in this drawing. Each row consists of p copies of the same vertex of σ^n, one for each factor in the deleted join, and since the join is 2-wise deleted, a simplex can use only one of the copies in each row.

Alternatively, we can also consider the maximal simplices as functions $[n+1] \to [p]$.

Notes. The original proof of Tverberg's theorem [Tve66] is complicated. The idea is simple, though: start with some point configuration for which the theorem is valid, and convert it to a given configuration by moving one point at a time. During the movement, the current partition may stop working at some point, and it must be shown that it can be replaced by another suitable partition. Later on, Tverberg found a simpler proof [Tve81]. Sarkaria [Sar92] invented another, very nice, and reasonably simple proof, based on a geometric lemma due to Bárány, and his proof was further streamlined by Onn (see [BO97]). Still another proof, also due to Tverberg and inspired by Bárány's proof, was published in Tverberg and Vrećica [TV93]. A similar proof, technically somewhat simpler, is due to Roudneff [Rou01].

Tverberg's theorem is quite important and has numerous applications, as well as extensions and generalizations; see e.g. Eckhoff [Eck93]. Some interesting aspects are briefly discussed in Kalai's lively survey [Kal01].

The topological Tverberg theorem. Bárány et al. [BSS81] proved Theorem 6.5.2 using deleted products. By an ingenious argument, they showed that the p-fold 2-wise deleted product of σ^N is $(N-p)$-connected. Then they established and used apparently the first theorem of Borsuk–Ulam type dealing with $\mathbb{Z}_p$-actions that has appeared in a combinatorial-geometric application. In our terminology, that result can be phrased as follows. For a prime p and integer $d \geq 1$, consider the sphere $S^{d(p-1)-1}$ with the $\mathbb{Z}_p$-action obtained as in the proof of Proposition 6.3.2, and let $X_{d,p} := S^{d(p-1)-1} * [p]$ (the action on $[p]$ is a cyclic permutation of the p points). Then for any continuous $f\colon X_{d,p} \to \mathbb{R}^d$ there is a point $x \in X_{d,p}$ whose whole orbit under the $\mathbb{Z}_p$-action is mapped to a single point in $\mathbb{R}^d$ (this is a special case of Theorem 6.3.3, and, in fact, it is equivalent to it). Their proof proceeds by reduction to a lemma from Krasnosel'skiĭ and Zabrejko [KZ75], claiming that, given a free $\mathbb{Z}_p$-action on S^n, any $\mathbb{Z}_p$-map $S^n \to S^n$ has degree 1 modulo p.

The technique of deleted joins for such problems was developed by Sarkaria [Sar90], [Sar91a].

The validity of the topological Tverberg theorem for arbitrary (non-prime) p is one of the most challenging problems in this field. For p being a prime power, the theorem was proved by Özaydin [Öza] in an unpublished manuscript, and much later by Volovikov [Vol96] (and also by Sarkaria [Sar00]). Assuming the theorem of Volovikov mentioned in the notes to Section 6.2 about maps of fixed-point free $(\mathbb{Z}_p \times \cdots \times \mathbb{Z}_p)$-spaces, the proof is a relatively straightforward generalization of the proof in this section.

Exercises

1. Prove (directly, without using Tverberg's Theorem) that for any integers $d, r_1, r_2 \geq 2$, we have $T(d, r_1 r_2) \leq T(d, r_1)T(d, r_2)$. (Together with Radon's theorem, this implies that $T(d,r)$ is finite for all d and r.)

2. Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d+1}$ be vertices of a simplex in $\mathbb{R}^d$ and let B_i be a set of $r-1$ points lying very close to $\boldsymbol{v}_i$. Prove that there is no partition of $B := B_1 \cup \cdots \cup B_{d+1}$ into r disjoint parts whose convex hulls have a nonempty intersection.

6.6 Many Tverberg partitions

A conjecture of Sierksma, still unresolved at the time of writing, states that the number of Tverberg partitions for a set of $(r-1)(d+1)+1$ points in $\mathbb{R}^d$ in general position is at least $((r-1)!)^d$. This number is attained for the configuration of $d+1$ tight clusters, with $r-1$ points each, placed at the vertices of a simplex, and one point in the middle (Exercise 1). (We count *unordered* partitions, where the order of the sets $A_1, \ldots, A_r$ does not matter.)

For a prime r, one can prove a quite good lower bound by cleverly extending the topological proof (while no non-topological method is known to yield a good lower bound).

6.6.1 Theorem (Many Tverberg partitions; Vućić and Živaljević [VŽ93]). *Let p be a prime. For any continuous map $f\colon \|\sigma^N\| \to \mathbb{R}^d$, where $N = (d+1)(p-1)$, the number of unordered p-tuples $\{F_1, F_2, \ldots, F_p\}$ of pairwise disjoint faces of σ^N with $\bigcap_{i=1}^p f(\|F_i\|) \neq \emptyset$ is at least*

$$\frac{1}{(p-1)!} \cdot \left(\frac{p}{2}\right)^{(d+1)(p-1)/2}.$$

We note that for d and p large, this bound is roughly the square root of the bound conjectured by Sierksma.

Proof. Let K denote the simplicial complex $(\sigma^N)^{*p}_{\Delta(2)}$. As we know, the maximal simplices of K are the edges of the complete $(N+1)$-partite hypergraph; if the vertex set of K is identified with $[N+1] \times [p]$, then such a maximal simplex S is $\{(1, i_1), (2, i_2), \ldots, (N+1, i_{N+1})\}$, $i_1, \ldots, i_{N+1} \in [p]$. Such an S encodes the *ordered* partition $(F_1, F_2, \ldots, F_p)$ with $F_i = \{j \in [N+1] : i_j = i\}$. For example, with $d = 2$ and $p = 3$, the indicated S in the picture encodes the ordered Tverberg partition of the $N+1 = 7$ points drawn on the right:

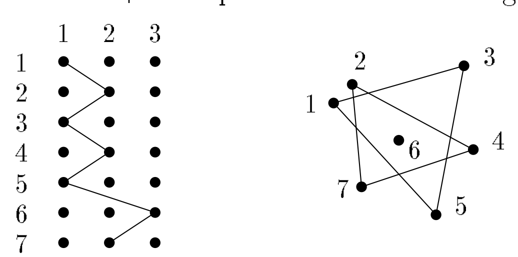

Call S *good* whenever it encodes a Tverberg partition; that is, whenever $\bigcap_{i=1}^{p} f(\|F_i\|) \neq \emptyset$. An S is good exactly if it contains a point mapped to the diagonal in $(\mathbb{R}^d)^{*p}$ by f^{*p}, where f^{*p} is the p-fold join of f as in the preceding proof. If we prove that K has at least M good maximal simplices, we obtain that there are at least $M/p!$ (unordered) Tverberg partitions.

Here is the strategy of further progress. We define a suitable family $\mathcal{L}$ of subcomplexes $\mathsf{L} \subset \mathsf{K}$. Each L is closed under the $\mathbb{Z}_p$-action (cyclic shift on the rows of the hypergraph), and $\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{L}) \geq N$ so that f^{*p} restricted to $\|\mathsf{L}\|$ maps some point to the diagonal. Consequently, each $\mathsf{L} \in \mathcal{L}$ contains a good maximal simplex (actually at least p of them). Finally, we count the number Q of $\mathsf{L} \in \mathcal{L}$ containing any given maximal simplex of K, and estimate $M \geq p \cdot |\mathcal{L}|/Q$.

Since in the case $p = 2$ the theorem is already proved, we may now assume $p > 2$, so p is odd, and $N = (p{+}1)(d{-}1)$ is even. To describe a member L of the family $\mathcal{L}$, we first divide the $N{+}1$ rows in the hypergraph into $\frac{N}{2}$ pairs plus one remaining row; let Π be the number of ways of accomplishing this (we do not need its value since it will cancel out later). Next, we look at the two rows in one of the pairs; the simplices of K living on these rows are the edges of the complete bipartite graph between the rows. We choose a cycle C in this complete bipartite graph that is invariant under the cyclic shift action. Some thought reveals that such a cycle is uniquely determined by choosing two distinct edges emanating from the first vertex of the top row into the bottom row, as in the drawing (for $p = 5$):

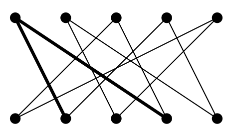

All the other edges are given as shifts of the chosen two. (Yes, we always get just one cycle—right?) Thus, there are $\binom{p}{2}$ choices for C. Such a cycle is chosen for each pair of rows, so we obtain invariant cycles $C_1, \ldots, C_{N/2}$. For a fixed pairing of the rows, the number of choices of the C_i is $\binom{p}{2}^{N/2}$. The maximal simplices of the subcomplex L corresponding to a given choice of the row pairing and of the C_i are the maximal simplices of K that contain an edge of each C_i, such as is drawn below:

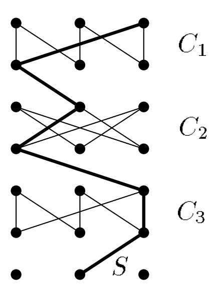

We have $|\mathcal{L}| = \Pi \cdot \binom{p}{2}^{N/2}$. We leave it as an exercise to show that the number Q of complexes $\mathsf{L} \in \mathcal{L}$ that contain a given maximal simplex $S \in \mathsf{K}$ is $\Pi \cdot (p{-}1)^{N/2}$.

Each L can be interpreted as the join of its $N/2$ cycles $C_1, \ldots, C_{N/2}$ and of the remaining p points. Thus, topologically,

$$\|\mathsf{L}\| \cong (S^1)^{*(N/2)} * [p] \cong S^{N-1} * [p],$$

and so $\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{L}) \geq N$ as required. Theorem 6.6.1 follows by the calculation indicated above.

> **Notes.** The presented proof of the lower bound for the number of Tverberg partitions is a simplification of the argument of Vućić and Živaljević [VŽ93] (instead of the invariant subcomplexes L, they consider non-invariant cones over invariant spheres in K and use an argument about mapping degrees).

Exercises

1. Show that the number of (unordered) Tverberg r-partitions for the configuration described in the text ($d{+}1$ clusters by $r{-}1$ points near the vertices of a simplex in $\mathbb{R}^d$ and one point in the center of the simplex) equals $((r{-}1)!)^d$.

6.7 $\mathbb{Z}_p$-index, Kneser colorings, and p-fold points

In this part, we more or less repeat the considerations about index and Kneser colorings from Section 5.7 in a p-fold setting. No new ideas are needed; one just has to get the definitions right and verify that the proofs work. As a reward, we then prove quite quickly some theorems which were generally considered reasonably hard,

Considering the proof of the topological Tverberg theorem in Section 6.5 and replacing σ^N with an arbitrary simplicial complex K, we obtain

6.7.1 Proposition (Index and p-fold points). *Let p be a prime and let K be a simplicial complex such that*

$$\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{K}^{*p}_{\Delta(2)}) \geq (d{+}1)(p{-}1).$$

Then for any continuous mapping $f\colon \|\mathsf{K}\| \to \mathbb{R}^d$ there are points $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_p \in \|\mathsf{K}\|$ with pairwise disjoint supports such that $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = \cdots = f(\boldsymbol{x}_p)$.

By Sarkaria's inequality (Proposition 6.2.4(vi)), if K is a subcomplex of a larger simplicial complex J for which we know $\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{J}^{*p}_{\Delta(2)})$, we can estimate

$$\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{K}^{*p}_{\Delta(2)}) \geq \mathrm{ind}_{\mathbb{Z}_p}(\mathsf{J}^{*p}_{\Delta(2)}) - \mathrm{ind}_{\mathbb{Z}_p}(\Delta(\mathsf{J}^{*p}_{\Delta(2)} \setminus \mathsf{K}^{*p}_{\Delta(2)})) - 1.$$

So we want to bound above the $\mathbb{Z}_p$-index of $\Delta(\mathsf{J}^{*p}_{\Delta(2)} \setminus \mathsf{K}^{*p}_{\Delta(2)})$, and this can be done using Kneser-like colorings.

Kneser hypergraphs. Let $\mathcal{S}$ be a set system. Generalizing the notion of Kneser graph, we define the *Kneser r-hypergraph* $\mathrm{KG}_r(\mathcal{S})$: the vertex set is $\mathcal{S}$ and the edges are all r-tuples of pairwise disjoint sets as edges; that is,

$$\Big\{\{S_1, S_2, \ldots, S_r\} : S_1, \ldots, S_r \in \mathcal{S}, S_i \cap S_j = \varnothing \text{ for } 1 \le i < j \le r\Big\}.$$

We recall that a proper m-coloring of a hypergraph H is a mapping $c\colon V(H) \to [m]$ such that no edge of H is monochromatic. For the Kneser r-hypergraph $\mathrm{KG}_r(\mathcal{S})$, we color the sets in $\mathcal{S}$ and we want that no r pairwise disjoint sets get the same color. Or, phrased differently, we want a coloring of the vertices of the usual Kneser graph $\mathrm{KG}(\mathcal{S})$ such that no clique (complete subgraph) of size r is monochromatic.

The following lemma gives a whole family of bounds for the $\mathbb{Z}_p$-index:

6.7.2 Lemma (Index bound from coloring KG_r). *Let p be a prime, let K be an invariant subcomplex of a simplicial $\mathbb{Z}_p$-complex J, and let $\mathcal{S} := \mathrm{MIN}(\mathsf{J} \setminus \mathsf{K})$ be the system of the inclusion-minimal simplices in $\mathsf{J} \setminus \mathsf{K}$. Then, for any $r = 2, 3, \ldots, p$, we have*

$$\mathrm{ind}_{\mathbb{Z}_p}(\Delta(\mathsf{J}^{*p}_{\Delta(2)} \setminus \mathsf{K}^{*p}_{\Delta(2)})) \quad\le\quad (r-1) \cdot \chi(\mathrm{KG}_r(\mathcal{S})) - 1.$$

We will actually use only the case $r = p$ (which tends to give the strongest bound, although it need not always be the case; see Exercise 3.

We remark that the restriction on p being prime is only needed to guarantee that the deleted join $(\sigma^{m-1})^{*p}_{\Delta(r)}$ is a free $\mathbb{Z}_p$-space. For $r = 2$, for example, any group of permutations of $[p]$, including $\mathbb{Z}_p$ represented by the cyclic shift, acts freely even if p is composite.

Proof of Lemma 6.7.2. As in the proof of Lemma 5.7.1, we define the labeling h of the simplices in J by subsets of $[m]$:

$$h(F) = \{c(G) : G \in \mathcal{S}, G \subseteq F\}.$$

Note that simplices in K receive $\varnothing$ while those in $\mathsf{J} \setminus \mathsf{K}$ receive a nonempty set. For a simplex $F_1 * \cdots * F_p \in \mathsf{J}^{*p}_{\Delta(2)} \setminus \mathsf{K}^{*p}_{\Delta(2)}$, we put $g(F_1 * \cdots * F_p) := h(F_1) * \cdots * h(F_p)$. Since c is a proper coloring of $\mathrm{KG}_r(\mathcal{S})$, each r sets among $h(F_1), \ldots, h(F_p)$ have an empty intersection, and so g is a simplicial $\mathbb{Z}_p$-map into $\mathrm{sd}((\sigma^{m-1})^{*p}_{\Delta(r)})$.

It remains to show that the index of the latter space is (at most) $m(r-1)-1$. This is left as Exercise 1.

Together with Proposition 6.7.1, Lemma 6.7.2 yields

6.7.3 Theorem (Sarkaria's theorem on coloring and p-fold points). *Let p be a prime. Let K be a subcomplex of a simplicial complex J, and suppose that for some $r \in \{2, 3, \ldots, p\}$,*

$$d \le \frac{1}{p-1}\,\mathrm{ind}_{\mathbb{Z}_p}(\mathsf{J}^{*p}_{\Delta(2)}) - \frac{r-1}{p-1}\,\chi(\mathrm{KG}_r(\mathrm{MIN}(\mathsf{J} \setminus \mathsf{K}))) - 1.$$

Then for any continuous map $f\colon \|\mathsf{K}\| \to \mathbb{R}^d$ there are p points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_p \in \|\mathsf{K}\|$ with pairwise disjoint supports such that $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = \cdots = f(\boldsymbol{x}_p)$. For $\mathsf{J} = \sigma^n$ and $r = p$, the condition is

$$d \le \frac{n}{p-1} - \chi(\mathrm{KG}_p(\mathcal{S})) - 1.$$

6.7.4 Example (Tverberg's theorem with restricted dimensions). In Tverberg's theorem, $(d+1)(r-1)-1$ points in $\mathbb{R}^d$ suffice to get r disjoint subsets with intersecting convex hulls. What happens if we consider $N+1$ points and want r disjoint subsets with intersecting convex hulls, but each of the sets should have at most $k+1 \le d$ points? For example, for $r = 3$, $d = 3$, and $k = 2$, we would like to find 3 vertex-disjoint triangles in $\mathbb{R}^3$ with a common point.

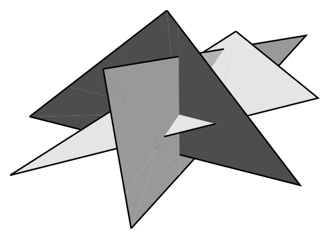

It turns out that such triangles always exist, even with the smallest conceivable number of points, i.e. 9. On the other hand, no matter how many points in suitable general position in $\mathbb{R}^3$ we have, we cannot find 4 vertex-disjoint intersecting triangles. More generally, if the sum of codimensions of the r convex hulls, i. e. $r(d-k)$, is greater than d, no N will do.

For r being a prime and such that the codimension condition $r(d-k) \le d$ holds, one can show the existence of a suitable N using Theorem 6.7.3; see Exercise 4.

Notes. This section is again based on [Sar91a] and [Sar90].

Sarkaria also considers the k-wise deleted joins $\mathsf{K}^{*p}_{\Delta(k)}$ (with K being the $(k-1)$-skeleton of an n-simplex) and uses Kneser-like colorings for determining the index of such deleted joins.

Example 6.7.4 is inspired by Vrećica and Živaljević [ŽV94].

*** Tverberg–Vrećica project, some results?

Kneser's conjecture. Here we briefly summarize refences concerning the chromatic number of Kneser graphs and hypergraphs. As was mentioned in Section 3.3, Kneser's conjecture [Kne55] was first proved by Lovász [Lov78]. (Previously Garey & Johnson [GJ76] had established the case $k = 3$ by elementary means; also see Stahl [Sta76].)

Lovász' proof is not included in our text (it may appear in a future version). With every graph G, Lovász associated a simplicial complex

$\mathsf{N}(G)$, whose vertex set is $V = V(G)$ and the simplices are the subsets vertices having a common neighbor in G. He then proved that if $\mathsf{N}(G)$ is k-connected, then $\chi(G) \geq k+3$, and analyzed the connectivity of the neighborhood complex of the Kneser graph. This approach was further developed in Alon, Frankl & Lovász [AFL86] (who generalized the results to hypergraphs, and defined the technically somewhat more convenient *box complex* associated with a graph or hypergraph; also see Kříž [Kri92]). Walker [Wal83] showed that graph homomorphisms induce $\mathbb{Z}_2$-maps of suitably modified neighborhood complexes (according to Björner [Bjö95], it was also independently noted by Lovász in unpublished lecture notes). For another application of neighborhood complexes see Lovász [Lov83].

Alon, Frankl & Lovász [AFL86] established Erdős' generalization of Kneser's conjecture for hypergraphs: if $n \geq (m-1)(r-1)+rk$, then $\chi(\mathrm{KG}_r(\binom{[n]}{k})) > m$.

By the method shown in Section 3.4, Dol'nikov estimated $\chi(\mathrm{KG}_r(\mathcal{S}))$ from below by the minimum cardinality of a set $Y \subseteq X$ (where X is the ground set of $\mathcal{S}$) such that $X \setminus Y$ can be colored by two colors so that no color class contains $\lceil \frac{r}{2} \rceil$ pairwise disjoint sets of $\mathcal{S}$. He then re-proves result of [AFL86] on $\chi(\mathrm{KG}_r(\binom{[n]}{k}))$ for all even r; for odd r he needs an additional condition on the parameters r, k, n. Yet another proof of a statement generalizing the Erdős' conjecture was given by Sarkaria [Sar90]; see Exercise 2.

Kříž [Kri92, Kri00] proved the following generalization of Dol'nikov's theorem: for any set system $\mathcal{S}$,

$$\chi(\mathrm{KG}_r(\mathcal{S})) \geq \frac{1}{r-1} \cdot \mathrm{cd}_r(\mathcal{S}),$$

where $\mathrm{cd}_r(\mathcal{S})$ is the r-colorability defect introduced in Section 3.4. This theorem, too, easily implies the results of Alon et al. on $\chi(\mathrm{KG}_r(\binom{[n]}{k}))$. The proof in [Kri92] does not work in the generality stated there (as was pointed out by Živaljević) but the result for the Kneser hypergraphs remains valid [Kri00]. A simplified version of Kříž's proof, emplyoing a Sarkaria-style inequality for estimating the index of a certain space, was given in [Mat01b] (see Exercise 5.7.3 for the special case of this proof with $r = 2$).

In [Mat01a], using the ideas from the just mentioned proof, Kneser's conjecture was derived from Tucker's lemma by a direct combinatorial argument, without using a continuous result of Borsuk–Ulam type. Since the required instance of Tucker's lemma also has a combinatorial proof, the resulting proof of Kneser's conjecture is purely combinatorial, although the topological inspiration remains notable, of course. An extensive generalization of this method was obtained by Ziegler [Zie01]. He formulated a $\mathbb{Z}_p$-analogue of the required special instance of Tucker's lemma, and derived many generalizations of Kneser's conjecture from it (including Schrijver's theorem, the Dol'nikov–Kříž theorem, and Sarkaria's results).

Exercises

1. (a) Prove that the $\mathbb{Z}_p$-index of the p-fold k-wise deleted join $(\sigma^n)^{*p}_{\Delta(k)}$ is at most $(n{+}1)(k{-}1){-}1$, $2 \leq k \leq p$.

   (b) Show that the index in (a) is actually equal to $(n{+}1)(k{-}1){-}1$.

2. (a) Find a coloring of the Kneser r-hypergraph $\mathrm{KG}_r(\binom{[n]}{k})$ by $\lceil \frac{n-r(k-1)}{r-1} \rceil$ colors.

   (b) Use Theorem 6.7.3 to prove that this number of colors is the smallest possible.

3. (a) Prove that for $r \geq 3$ and any finite set system $\mathcal{S}$, we have

$$\chi(\mathrm{KG}_r(\mathcal{S})) \leq \left\lceil \frac{\chi(\mathrm{KG}_2(\mathcal{S}))}{r-1} \right\rceil.$$

   (b) More generally, check that for $r > q \geq 2$, we have

$$\chi(\mathrm{KG}_r(\mathcal{S})) \leq \left\lceil \frac{\chi(\mathrm{KG}_q(\mathcal{S}))}{\frac{r}{q-1} - 1} \right\rceil.$$

4. (a) Let p be a prime and let f be a continuous map of the k-skeleton of the N-simplex into $\mathbb{R}^d$. Supposing that $p(d{-}k) \leq d$ and $N = (d{+}2)(p{-}1)$, use Theorem 6.7.3 to show that there are p points with pairwise disjoint supports that are mapped to the same point by f.

   (b) Derive the claim of Example 6.7.4 from (a): for any prime p, any $d \geq 1$, and any k, $0 \leq k \leq d{-}1$, such that $p(d{-}k) \leq d$, there exists N such that among any N points in $\mathbb{R}^d$, one can select p disjoint groups of size $k{+}1$ each whose convex hulls have a nonempty intersection.

   (c) Show that the conclusion of (b) need not be true for any N if the codimension condition $p(d{-}k) \leq d$ is not satisfied.

6.8 The colored Tverberg theorem

If we have 7 points in the plane, by Tverberg's theorem we can divide them into 3 groups whose convex hulls have a common intersection. The colored version of this statement is: given 3 red, 3 blue, and 3 white points in the plane, we can always partition them into 3 "tricolores" with intersecting convex hulls:

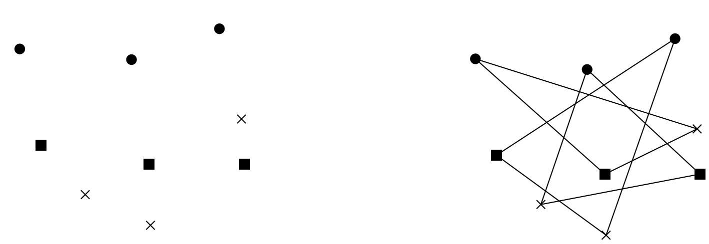

6.8.1 Theorem (The colored Tverberg theorem). *For any integers $r \geq 2$ and $d \geq 1$ there exists an integer $t = t(d, r)$, such that for any $d{+}1$ pairwise disjoint t-point sets $C_1, C_2, \ldots, C_{d+1}$, we can find disjoint sets $A_1, A_2, \ldots, A_r$ with $|A_i \cap C_j| = 1$ for all $i = 1, 2, \ldots, r$ and $j = 1, 2, \ldots, d{+}1$ such that $\bigcap_{i=1}^{r} \mathrm{conv}(A_i) \neq \varnothing$. If we think of the points of C_j as having color j, then each A_i is required to use all colors (to be a "rainbow" set).*

This may look like an innocent (and not too exciting) variation of Tverberg's theorem, but in fact, this theorem attracted great interest. It was a key ingredient in obtaining a nontrivial upper bound in the so-called *k-set problem*: what is the maximum number of distinct k-element subsets of an n-point set $A \subset \mathbb{R}^d$ that can be cut off by a halfspace, i. e. what is

$$\max_{A \subset \mathbb{R}^d, |A|=n} |\{A \cap h : |A \cap h| = k, \ \ h \text{ a halfspace}\}|.$$

This problem seems to be very hard even in the plane. Here we will not explain the connection to the colored Tverberg theorem (see, e.g., [ABFK92]).

While the Tverberg theorem can be proved in an elementary way, all known proofs for the colored version are topological.

We prove the following topological version, which implies the colored Tverberg theorem with $t = 4r{-}1$. We use "Bertrand's postulate," which states that for any $r > 1$ there is a prime p with $r \leq p < 2r$. (This was first proved by Chebyshev, and Erdős found the first simple and elementary proof as a first-year undergraduate at age 17 [Erd32] [AZ00].)

6.8.2 Theorem (Topological colored Tverberg theorem [ŽV92]). *Let d be a positive integer and let p be a prime. Let $C_1, C_2, \ldots, C_{d+1}$ be disjoint sets of cardinality $2p{-}1$ each, and let K be the simplicial complex with vertex set $C_1 \dot\cup C_2 \dot\cup \cdots \dot\cup C_{d+1}$, whose simplices are all subsets using at most one point from each C_i. (In other words, $\mathsf{K} = [2p{-}1]^{*(d+1)}$.) Then for any continuous map $f\colon \|\mathsf{K}\| \to \mathbb{R}^d$, there are p pairwise disjoint faces $F_1, F_2, \ldots, F_p$ of K whose images intersect: $\bigcap_{i=1}^{p} f(\|F_i\|) \neq \varnothing$.*

Proof. With the powerful Theorem 6.7.3 on coloring and p-fold points, the proof is routine.

We take $\mathsf{J} := \sigma^N$, where $N = |V(\mathsf{K})| - 1 = (d{+}1)(2p{-}1){-}1$. The system $\mathcal{S}$ of minimal non-faces of K consists of all edges connecting two points in the same C_i. We work with $r = p$, i. e. we look for a coloring of the Kneser p-hypergraph $\mathrm{KG}_p(\mathcal{S})$. Having p disjoint edges of $\mathcal{S}$, they together cover $2p$ points, and so they cannot all live on the same class C_i. Thus, coloring all the edges on C_i by the color i shows $\chi(\mathrm{KG}_p(\mathcal{S})) \leq d{+}1$. The right-hand side in the condition in Theorem 6.7.3 comes out as $d \cdot \frac{p}{p-1} > d$, and we are done.

Notes. The colored Tverberg theorem was proved for $d = 2$ and conjectured to hold for arbitrary d by Bárány, Füredi & Lovász [BFL90]. The general d-dimensional case was proved, with $t \leq 4r{-}1$, by Živaljević & Vrećica [ŽV92]. A simpler proof was found by Björner, Lovász, Vrećica & Živaljević [BLŽV94] (also see [ŽV94] for a similar argument and Ziegler [Zie94] for yet another approach). The proof of the colored Tverberg theorem by Sarkaria's method was noted in [Mat96].

Bárány et al. actually conjectured that $t = r$ should suffice in the colored Tverberg theorem. This is known for $d = 2$ (Bárány & Larman [BL92]) and for $r = 2$ (Lovász; also published in [BL92]—the Borsuk–Ulam theorem is applied in a beautiful way). For r a prime, the Živaljević–Vrećica approach gives $t \leq 2r-1$. This was extended to all prime powers r by Živaljević [Živ98] (similar to the proofs of the topological Tverberg theorem for prime powers, as mentioned in the notes to Section 6.5).

Exercises

1. (Vrećica & Živaljević [ŽV94]) This is a colored version of Example 6.7.4.

   (a) Given 5 red, 5 blue, and 5 red points in $\mathbb{R}^3$, prove that there are 3 vertex-disjoint tricolored triangles having a common point.

   (b) Let $C_1, \ldots, C_{k+1} \subset \mathbb{R}^d$ be sets of cardinality $2p-1$ each, where p is a prime satisfying $p(d-k) \leq d$. Prove that there are p pairwise disjoint rainbow sets $A_1, \ldots, A_p$ (with $|A_i \cap C_j| = 1$ for all i, j) such that $\bigcap_{i=1}^{p} \operatorname{conv}(A_i) \neq \emptyset$.

Bibliography

[AA89] J. Akiyama and N. Alon. Disjoint simplices and geometric hypergraphs. In *Combinatorial Mathematics*, volume 555 of *Annals New York Acad. Sci.*, pages 1–3, 1989. Proc. 3rd Int. Conf., New York 1985.

[ABFK92] N. Alon, I. Bárány, Z. Füredi, and D. Kleitman. Point selections and weak ε-nets for convex hulls. *Combinatorics, Probabability Comput.*, 1:189–200, 1992.

[AE98] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In *Discrete and Computational Geometry: Ten Years Later (B. Chazelle, E. Goodman, and R. Pollack eds.)*. American Mathematical Society, Providence, 1998. In press.

[AFL86] N. Alon, P. Frankl, and L. Lovász. The chromatic number of Kneser hypergraphs. *Transactions Amer. Math. Soc.*, 298:359–370, 1986.

[AKK$^+$00] J. Akiyama, A. Kaneko, M. Kano, G. Nakamura, E. Rivera-Campo, S. Tokunaga, and J. Urrutia. Radial perfect partitions of convex sets in the plane. In *Discrete and Computational Geometry (J. Akiyama et al. eds), Lect. Notes Comput. Sci. 1763*, pages 1–13. Springer, Berlin, 2000.

[Alo87] N. Alon. Splitting necklaces. *Advances in Math.*, 63:247–253, 1987.

[Alo88] N. Alon. Some recent combinatorial applications of Borsuk-type theorems. In *Algebraic, Extremal and Metric Combinatorics*, volume 131 of *Lond. Math. Soc. Lecture Note Series*, pages 1–12, 1988. Pap. Conf., Montreal 1986.

[AP83] J. K. Arason and A. Pfister. Quadratische Formen über affinen Algebren und ein algebraisches Beweis des Satzes von Borsuk–Ulam. *J. Reine Angew. Math.*, 339:163–164, 1983.

[Avi85] D. Avis. On the partitionability of point sets in space. In *Proc. 1st ACM Symposium on Computational Geometry*, pages 116–120, 1985.

[AW86] N. Alon and D. B. West. The Borsuk-Ulam theorem and bisection of necklaces. *Proc. Amer. Math. Soc.*, 98:623–628, 1986.

[AZ00] M. Aigner and G. M. Ziegler. *Proofs from THE BOOK, 2nd edition.* Springer-Verlag, 2000.

[Bár78] I. Bárány. A short proof of Kneser's conjecture. *J. Combinatorial Theory, Ser. A*, 25:325–326, 1978.

[Bár80] I. Bárány. Borsuk's theorem through complementary pivoting. *Math. Programming*, 18:84–88, 1980.

[Bár93] I. Bárány. Geometric and combinatorial applications of Borsuk's theorem. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, chapter IX, pages 235–249. Springer-Verlag, Berlin Heidelberg, 1993.

[BB79] E. G. Bajmóczy and I. Bárány. A common generalization of Borsuk's and Radon's theorem. *Acta Math. Hungarica*, 34:347–350, 1979.

[BB89] I. K. Babenko and S. A. Bogatyj. Mapping a sphere into euclidean space. *Math. Notes*, 46(3):683–686, 1989. translation from *Mat. Zametki* 46, No. 3 (1989), 3-8.

[BFL90] I. Bárány, Z. Füredi, and L. Lovász. On the number of halving planes. *Combinatorica*, 10:175–183, 1990.

[Bie] T. Bier. A remark on alexanderduality and the disjunct join. Preprint 1992, 8 pages.

[Bjö95] A. Björner. Topological methods. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume II, chapter 34, pages 1819–1872. North-Holland, Amsterdam, 1995.

[BKS00] S. Bespamyatnikh, D. Kirkpatrick, and J. Snoeyink. Generalizing ham sandwich cuts to equitable subdivisions. *Discrete Comput. Geom.*, 24:605–622, 2000.

[BL92] I. Bárány and D. Larman. A colored version of Tverberg's theorem. *J. London Math. Soc. II. Ser.*, 45:314–320, 1992.

[BLŽV94] A. Björner, L. Lovász, R. Živaljević, and S. Vrećica. Chessboard complexes and matching complexes. *J. London Math. Soc.*, 49:25–39, 1994.

[BM] I. Bárány and J. Matoušek. Equipartition of two measures by a 4-fan. *Discrete Comput. Geom.* To appear.

[BM01] I. Bárány and J. Matoušek. Simultaneous partitions of measures by k-fans. *Discrete Comput. Geom.*, 25:317–334, 2001.

[BO97] I. Bárány and S. Onn. Colourful linear programming and its relatives. *Math. Oper. Res.*, 22:550–567, 1997.

[Bor33] K. Borsuk. Drei Sätze über die n-dimensionale euklidische Sphäre. *Fundamenta Mathematicae*, 20:177–190, 1933.

[Bou55] D. G. Bourgin. On some separation and mapping theorems. *Comment. Math. Helv.*, 29:199–214, 1955.

[Bou63] D. G. Bourgin. *Modern Algebraic Topology*. The Macmillan Company, New York, 1963.

[Bre72] G. E. Bredon. *Introduction to Compact Transformation Groups*. Academic Press, New York, 1972.

[Bre93] G. Bredon. *Topology and Geometry (Graduate Texts in Mathematics 139)*. Springer-Verlag, Berlin etc., 1993.

[BS92] U. Brehm and K. Sarkaria. Linear vs. piecewise linear embeddability of simplicial complexes. Tech. Report 92/52, Max-Planck-Institut f. Mathematik, Bonn, Germany, 1992.

[BSS81] I. Bárány, S. B. Shlosman, and S. Szücs. On a topological generalization of a theorem of Tverberg. *J. London Math. Soc., II. Ser.*, 23:158–164, 1981.

[BT93] G. Brightwell and W. T. Trotter. The order dimension of convex polytopes. *SIAM J. Discrete Math.*, 6(2):230–245, 1993.

[BT97] G. Brightwell and W. T. Trotter. The order dimension of planar maps. *SIAM J. Discrete Math.*, 10(4):515–528, 1997.

[BW95] A. Björner and V. Welker. The homology of "k-equal" manifolds and related partition lattices. *Advances in Math.*, 110:277–313, 1995.

[CHW92] G. Chartrand, H. Hevia, and R. J. Wilson. The ubiquitous Petersen graph. *"The Julius Petersen graph theory centennial," Discrete Math.*, 100:303–311, 1992.

[Col99] Y. Colin de Verdière. *Spectres de Graphes*, volume 4 of *Cours Spécialisés*. Société Mathématique de France, 1999.

[CW85] G. Chartrand and R. J. Wilson. The Petersen graph. In *"Graphs and Applications," Proc. 1st Symp. Graph Theory (Boulder 1982)*, pages 69–100. Wiley, 1985.

[DE94] T. K. Dey and H. Edelsbrunner. Counting triangle crossings and halving planes. *Discrete Comput. Geom.*, 12:281–289, 1994.

[dL99] M. de Longueville. Notes on the topological Tverberg theorem. Manuscript, TU Berlin, 1999.

[Dol'81] V.L. Dol'nikov. Transversals of families of sets. In *Studies in the theory of functions of several real variables (Russian)*, pages 30–36, 109. Yaroslav. Gos. Univ., Yaroslavl', 1981.

[Dol'92] V. L. Dol'nikov. A generalization of the ham sandwich theorem. *Math. Notes*, 52:771–779, 1992.

[Dol'94] V. L. Dol'nikov. Transversals of families of sets in R^n and a connection between the Helly and Borsuk theorems. *Russian Acad. Sci. Sb. Math.*, 79:93–107, 1994. Translated from *Ross. Akad. Nauk Matem. Sbornik,* Tom 184, No. 5, 1993.

[Dol83] A. Dold. Simple proofs of some Borsuk-Ulam results. *Contemp. Math.*, 19:65–69, 1983.

[DP97] C. T. J. Dodson and P. E. Parker. *User's Guide to Algebraic Topology*, volume 387 of *Mathematics and Its Applications.* Kluwer, Dordrecht, Boston, London, 1997.

[DP98] T. Day and J. Pach. Extremal problems for geometric hypergraphs. *Discr. Comput. Geom.*, 19:473–484, 1998.

[Dys51] F. J. Dyson. Continuous functions defined on spheres. *Ann. of Math. (2)*, 54:534–536, 1951.

[Eck79] J. Eckhoff. Radon's theorem revisited. In J. Tölke and J. Wills, editors, *Contributions to Geometry, Proc. Geometry Symposium, Siegen 1978*, pages 164–185. Birkhäuser, Basel, 1979.

[Eck93] J. Eckhoff. Helly, Radon and Carathéodory type theorems. In *Handbook of Convex Geometry (P. Gruber and J. Wills eds.)*, pages 389–448. North-Holland, Amsterdam, 1993.

[Ede87] H. Edelsbrunner. *Algorithms in Computational Geometry*, volume 10 of *EATCS Monographs in Theoretical Computer Science.* Springer-Verlag, Berlin, 1987.

[Eil40] S. Eilenberg. On a theorem of P. A. Smith concerning fixed points for periodic transformations. *Duke Math. J.*, 6:428–437, 1940.

[Eng77] R. Engelking. *General topology.* PNW Warszawa, 1977.

[Erd32] P. Erdős. Beweis eines Satzes von Tschebyschef. *Acta Math. Szeged*, 5:194–198, 1930/32.

[FH88] E. Fadell and S. Husseini. An ideal-valued cohomological index theory with applications to Borsuk–Ulam and Bourgin–Yang theorems. *Ergodic Theory Dyn. Syst. (Charles Conley Mem. Vol.)*, 8:73–85, 1988.

[Flo34] A. Flores. Über n-dimensionale Komplexe, die im R_{2n+1} absolut selbstverschlungen sind. *Ergeb. Math. Kolloq.*, 6:4–7, 1932/1934.

[FT81] R. M. Freund and M. J. Todd. A constructive proof of Tucker's combinatorial lemma. *J. Combinatorial Theory, Ser. A*, 30:321–325, 1981.

[Gal56] D. Gale. Neighboring vertices on a convex polyhedron. In H. W. Kuhn and A. W. Tucker, editors, *Linear Inequalities and Related Systems*, volume 38 of *Annals of Math. Studies*, pages 255–263, Princeton, 1956. Princeton University Press.

[GJ76] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *J. Association Computing Machinery*, 23:43–49, 1976.

[GP86] J. E. Goodman and R. Pollack. Upper bounds for configurations and polytopes in r^d. *Discrete Comput. Geometry*, 1:219–227, 1986.

[Grü67] B. Grünbaum. *Convex Polytopes.* Interscience, London, 1967.

[GW85] C. H. Goldberg and D. West. Bisection of circle colorings. *SIAM J. Algebraic Discrete Methods*, 6(1):93–106, 1985.

[Hat01] A. Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2001. In press. Preliminary version available at `http://math.cornell.edu/~hatcher#AT1`.

[Hir37] G. Hirsch. Une géneéralisation d'un theéorème de M. Borsuk concernant certaines transformations de l'*analysis situs*. *Acad. Roy. Belgique Bull. Cl. Sci.*, 23:219–225, 1937.

[Hir43] G. Hirsch. Sur des propriétés de représentations permutables et des généralisations d'un théorème de Borsuk. *Ann. Sci. Ecole Norm. Sup.*, 60(3):113–142, 1943.

[HR65] C. R. Hobby and J. R. Rice. A moment problem in L_1 approximation. *Proc. Amer. Math. Soc.*, 16:665–670, 1965.

[HS93] D. A. Holton and J. Sheehan. *The Petersen Graph*, volume 7 of *Australian Mathematical Society Lecture Series.* Cambridge University Press, Cambridge, 1993.

[IUY00] H. Ito, H. Uehara, and M. Yokoyama. 2-dimension ham-sandwich theorem for partitioning into three convex pieces. In *Discrete and Computational Geometry (J. Akiyama et al. eds), Lect. Notes Comput. Sci. 1763*, pages 129–157. Springer, Berlin, 2000.

[Kak43] S. Kakutani. A proof that there exists a circumscribing cube around any bounded closed convex set in R^3. *Ann. of Math. (2)*, 43:739–741, 1943.

[Kal88] G. Kalai. Many triangulated spheres. *Discrete Comput. Geometry*, 3:1–14, 1988.

[Kal01] G. Kalai. Combinatorics with a geometric flavor: Some examples. In *Visions in mathematics towards 2000 (GAFA, special volume), part II*, pages 742–792. Birkhäuser, Basel, 2001.

[KB83] W. Kühnel and T. F. Banchoff. The 9-vertex complex projective plane. *Math. Intelligencer*, 5:11–22, 1983.

[KK93] J. Kahn and G. Kalai. A counterexample to Borsuk's conjecture. *Bull. Am. Math. Soc., New Ser.*, 29:60–62, 1993.

[KK99] A. Kaneko and M. Kano. Balanced partitions of two sets of points in the plane. *Comput. Geom. Theor. Appl.*, 13(4):253–261, 1999.

[Kna47] B. Knaster. Problem 4. *Colloq. Math.*, 1:30, 1947.

[Kne55] M. Kneser. Aufgabe 360. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 58:2. Abteilung, S. 27, 1955.

[Kne82] M. Knebusch. An algebraic proof of the Borsuk–Ulam theorem for polynomial mappings. *Proc. Amer. Math. Soc.*, 84:29–32, 1982.

[Kra52] M. A. Krasnosel'skiĭ. On the estimation of the number of critical points of functionals (in Russian). *Uspehi Mat. Nauk*, 7:157–164, 1952.

[Kri92] I. Kriz. Equivariant cohomology and lower bounds for chromatic numbers. *Transactions Amer. Math. Soc.*, 33:567–577, 1992.

[Kri00] I. Kriz. A correction to "Equivariant cohomology and lower bounds for chromatic numbers". *Transactions Amer. Math. Soc.*, 352:1951–1952, 2000.

[Küh95] W. Kühnel. *Tight polyhedral submanifolds and tight triangulations*, volume 1612 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin Heidelberg, 1995.

[KZ75] M. A. Krasnoselsky and P. P. Zabrejko. *Geometrical methods of nonlinear analysis (in Russian)*. Nauka, Moscow, 1975. English translation: Springer, 1984.

[Lef49] S. Lefschetz. *Introduction to Topology*. Princeton, 1949.

[Lov78] L. Lovász. Kneser's conjecture, chromatic number and homotopy. *J. Combinatorial Theory, Ser. A*, 25:319–324, 1978.

[Lov83] L. Lovász. Self-dual polytopes and the chromatic number of distance graphs on the sphere. *Acta Sci. Math. (Szeged)*, 45:317–323, 1983.

[LS30] L. Lusternik and S. Schnirelmann. *Topological methods in variational problems*. Issledowatelskiĭ Institut Matematiki i Mechaniki pri O. M. G. U., Moscow, 1930. in Russian.

[LS98] L. Lovász and A. Schrijver. A Borsuk theorem for antipodal links and a spectral characterization of linklessly embeddable graphs. *Proc. Amer. Math. Soc.*, 126:275–1285, 1998.

[Mak84] V. V. Makeev. Spatial generalization of theorems on convex figures (in Russian). *Mat. Zametki*, 36(3):405–415, 1984. English translation: *Math. Notes* 36:700–705, 1984.

[Mak88] V. V. Makeev. Six-lobed partitions of three-dimensional space. *Vestn. Leningr. Univ.*, Ser. I, 1988(2):31–34, 1988. In Russian; English translation in *Vestn. Leningr. Univ.*, Math. 21, No.2, 40-45 (1988).

[Mak96] V. V. Makeev. Applications of topology to some problems in combinatorial geometry. In *Mathematics in St. Petersburg. Based on a three-day conference, St. Petersburg, Russia, October 1993 (A. A. Bolibruch ed.), Amer. Math. Soc. Transl., Ser. 2, 174*, pages 223–228. Amer. Math. Soc., Providence, Rhode Island, 1996.

[Mat95] J. Matoušek. Geometric range searching. *ACM Comput. Surveys*, 26:421–461, 1995.

[Mat96] J. Matoušek. Note on the colored Tverberg theorem. *J. Combin. Theory Ser. B*, 66:146–151, 1996.

[Mat01a] J. Matoušek. A combinatorial proof of Kneser's conjecture. *Combinatorica*, 2001. In press.

[Mat01b] J. Matoušek. On the chromatic number of Kneser hypergraphs. *Proc. Amer. Math. Soc.*, 2001. In press.

[Mun84] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, Reading, MA, 1984.

[Mun00] J. R. Munkres. *Topology. 2nd ed.* Prentice Hall, Upper Saddle River, NJ, 2000.

[Nil94] A. Nilli. On Borsuk's problem. In H. Barcelo and G. Kalai, editors, *Jerusalem Combinatorics '93*, volume 178 of *Contemporary Mathematics*, pages 209–210. Amer. Math. Soc., 1994.

[Nov00] I. Novik. A note on geometric embeddings of simplicial complexes in a Euclidean space. *Discrete Comput. Geom.*, 23(2):293–302, 2000.

[Oss99] P. Ossona de Mendez. Geometric realization of simplicial complexes. In *Graph Drawing 99 (J. Kratochíl ed.), Lect. Notes Comput. Sci. 1731*, pages 323–332. Springer, Berlin, 1999.

[Öza] M. Özaydin. Equivariant maps for the symmetric group. Preprint 1987, 17 pages.

[Pac86] U. Pachner. Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulationen semilinearer Mannigfaltigkeiten. *Abh. Math. Sem. Hamburg*, 57:69–86, 1986.

[Ram96] E. A. Ramos. Equipartition of mass distributions by hyperplanes. *Discrete Comput. Geom.*, 15:147–167, 1996.

[RC81] J. P. Robinson and M. Cohen. Counting sequences. *IEEE Trans. Comput.*, C-30:17–23, 1981.

[Rol90] D. Rolfsen. *Knots and links. 2nd print. with corr.* Mathematics Lecture Series 7, Publish or Perish, Houston, Texas, 1990.

[Rou01] J.-P. Roudneff. Partitions of points into simplices with k-dimensional intersection. Part I: The conic Tverberg's theorem. *European J. Combinatorics*, 22:733–743, 2001.

[Rud74] W. Rudin. *Real and Complex Analysis.* McGraw-Hill, New York, 1974.

[Sak] T. Sakai. Radial partitions of point sets in R^2. Manuscript, Tokoha Gakuen University, 1998.

[Sar90] K. S. Sarkaria. A generalized Kneser conjecture. *J. Combinatorial Theory, Ser. B*, 49:236–240, 1990.

[Sar91a] K. Sarkaria. A generalized van Kampen-Flores theorem. *Proc. Amer. Math. Soc.*, 111:559–565, 1991.

[Sar91b] K. S. Sarkaria. Kuratowski complexes. *Topology*, 30:67–76, 1991.

[Sar92] K. Sarkaria. Tverberg's theorem via number fields. *Isr. J. Math.*, 79:317, 1992.

[Sar00] K. S. Sarkaria. Tverberg partitions and Borsuk-Ulam theorems. *Pacific J. Math.*, 196:231–241, 2000.

[Sch78] A. Schrijver. Vertex-critical subgraphs of Kneser graphs. *Nieuw Arch. Wiskd., III. Ser.*, 26:454–461, 1978.

[Sch89] W. Schnyder. Planar graphs and poset dimension. *Order*, 5:323–343, 1989.

[Sch93] L. J. Schulman. An equipartition of planar sets. *Discrete Comput. Geometry*, 9:257–266, 1993.

[Sha57] A. Shapiro. Obstructions to the imbedding of a complex in a euclidean space. I: The first obstruction. *Ann. of Math., II. Ser.*, 66:256–269, 1957.

[Smi38] P. A. Smith. Transformations of finite period. *Annals of Math.*, 39:127–164, 1938.

[Smi41] P. A. Smith. Fixed point theorems for periodic transformations. *American J. Mathematics*, 63:1–8, 1941.

[Smi42] P. A. Smith. Fixed points of periodic transformations. In *Algebraic Topology, by S. Lefschetz*, volume 27 of *American Math. Soc. Colloq. Publications*, pages 350–373. Amer. Math. Soc., Providence RI, 1942.

[SS78] L. A. Steen and J. A. Seebach jun. *Counterexamples in topology. 2nd ed.* Springer-Verlag, New York, 1978.

[Sta76] S. Stahl. n-tuple colorings and associated graphs. *J. Combinatorial Theory, Ser. B*, 20:185–203, 1976.

[Ste85] H. Steinlein. Borsuk's antipodal theorem and its generalizations and applications: A survey. In A. Granas, editor, *Méthodes topologiques en analyse nonlinéaire*, volume 95 of *Colloqu. Sémin. Math. Super., Semin. Sci. OTAN (NATO Advanced Study Institute)*, pages 166–235, Montréal, 1985. Univ. de Montréal Press.

[Ste93] H. Steinlein. Spheres and symmetry: Borsuk's antipodal theorem. *Topol. Methods Nonlinear Anal.*, 1:15–33, 1993.

[Sti93] J. Stillwell. *Classical topology and combinatorial group theory. 2nd ed.* Graduate Texts in Mathematics 72. Springer-Verlag, New York, 1993.

[Šva57] A. S. Švarc. Some estimates of the genus of a topological space in the sense of Krasnosel'skiĭ (in Russian). *Uspehi Mat. Nauk*, 12(4):209–214, 1957.

[Šva62] A. S. Švarc. The genus of a fibre space (in Russian). *Trudy Moskov. Mat. Obšč.* 10:217–272, 1961 and, 11:99–126, 1962. English translation: *Amer. Math. Soc., Translat.*, II. Ser. 55:49–140, 1966.

[Tho92] C. Thomassen. The Jordan–Schoenflies theorem and the classification of surfaces. *Amer. Math. Monthly*, 99:116–130, 1992.

[Tuc46] A. W. Tucker. Some topological properties of disk and sphere. In *Proc. First Canadian Math. Congress (Montreal, 1945)*, pages 285–309, Toronto, 1946. University of Toronto Press.

[TV93] H. Tverberg and S. Vrečica. On generalizations of Radon's theorem and the Ham sandwich theorem. *European J. Comb.*, 14:259–264, 1993.

[Tve66] H. Tverberg. A generalization of Radon's theorem. *J. London Math. Soc.*, 41:123–128, 1966.

[Tve81] H. Tverberg. A generalization of Radon's theorem. II. *Bull. Aust. Math. Soc.*, 24:321–325, 1981.

[Vas92] V. A. Vassiliev. *Complements of discriminants of smooth maps: Topology and applications (Transl. of Math. Monographs 98).* Amer. Math. Soc., Providence, Rhode Island, 1992.

[vK32] R. E. van Kampen. Komplexe in euklidischen Räumen. *Abh. Math. Sem. Hamburg*, 9:72–78, 1932. Berichtigung dazu, *ibid.* (1932) 152–153.

[Vol96] A. Yu. Volovikov. On a topological generalization of the Tverberg theorem. *Math. Notes*, 59(3):324–326, 1996. Translation from Mat. Zametki 59, No.3, 454-456 (1996).

[VŽ93] A. Vućić and R. Živaljević. Note on a conjecture of Sierksma. *Discr. Comput. Geom*, 9:339–349, 1993.

[Wal83] J. W. Walker. From graphs to ortholattices to equivariant maps. *J. Comb. Theory Ser. B*, 35:171–192, 1983.

[Wu65] W.-T. Wu. *A theory of imbedding, immersion, and isotopy of polytopes in a Euclidean space.* Science Press, Peking, 1965.

[Yan54] C.-T. Yang. On theorems of Borsuk–Ulam, Kakutani–Yamabe-Yujobô and Dynson, I. *Annals of Math.*, 60:262–282, 1954.

[YY50] H. Yamabe and Z. Yujobô. On the continuous functions defined on a sphere. *Osaka Math. J.*, 2:19–22, 1950.

[Zie94] G. M. Ziegler. Shellability of chessboard complexes. *Israel J. Math.*, 87:97–110, 1994.

[Zie01] G. M. Ziegler. Generalized Kneser coloring theorems with combinatorial proofs. *Invent. Math.*, 2001. In press.

[Živ96] R. T. Živaljević. User's guide to equivariant methods in combinatorics. *Publ. Inst. Math. Beograde*, 59(73):114–130, 1996.

[Živ97] R. T. Živaljević. Topological methods. In J. E. Goodman and J. O'Rourke, editors, *CRC Handbook on Discrete and Computational Geometry*, pages 209–224, Boca Raton FL, 1997. CRC Press.

[Živ98] R. T. Živaljević. User's guide to equivariant methods in combinatorics. II. *Publ. Inst. Math. (Beograd) (N.S.)*, 64(78):107–132, 1998.

[ŽV90] R. T. Živaljević and S. T. Vrećica. An extension of the ham sandwich theorem. *Bull. London Math. Soc.*, 22:183–186, 1990.

[ŽV92] R. T. Živaljević and S. Vrećica. The colored Tverberg's problem and complexes of injective functions. *J. Combin. Theory, Ser. A*, 61:309–318, 1992.

[ŽV94] R. T. Živaljević and S. T. Vrećica. New cases of the colored Tverberg theorem. In *Jerusalem combinatorics '93: an international conference in combinatorics, May 9–17, 1993, Jerusalem, Israel (H. Barcelo ed.), Contemp. Math. 178*, pages 325–334. Amer. Math. Soc., Providence, Rhode Island, 1994.

[ŽV01] R. T. Živaljević and S. T. Vrećica. Conical equipartitons of mass distributions. *Discrete Comput. Geom.*, 25:335–350, 2001.

Index