We should hold day with the Antipodes,
If you would walk in absence of the sun.

William Shakespeare, *The Merchant of Venice*

This talk is almost entirely based off Jiří Matoušek's book *Using the Borsuk-Ulam Theorem.*

1 Preliminaries: The Borsuk-Ulam Theorem

The use of topology in combinatorics might seem a bit odd, but I would actually argue it has a long history. For instance, the existence of a Nash equilibrium is a famous quasi-combinatorial theorem whose only known proofs use topology in a crucial way.

The main tool we will use in this talk is the Borsuk-Ulam Theorem; here are several equivalent statements:

Theorem 1 (Borsuk 1933)**.**

1. *For every continuous map $f : S^n \to \mathbb{R}^n$, there is some $x \in S^n$ with $f(x) = f(-x)$.*
2. *For every antipodal (aka odd) mapping $g : S^n \to \mathbb{R}^n$, there is some $x \in S^n$ with $g(x) = 0$.*
3. *There is no antipodal map $h : S^n \to S^{n-1}$.*
4. *(Lyusternik and Shnirel'man 1930) For any covering $S^n = A_1 \cup \cdots \cup A_{n+1}$ such that each A_i is either open or closed, there is some A_i that contains an antipodal pair, namely there is some $x \in S^n$ with $x, -x \in A_i$.*

Proof. I won't prove the Borsuk-Ulam theorem. However, we can check that these statements are indeed equivalent.

(1) ⇒ (2) is immediate, since an antipodal map that agrees on $x, -x$ must map them both to 0.

(2) ⇒ (1) follows by defining $g(x) = f(x) - f(-x)$.

(2) ⇒ (3) is immediate, since there is an embedding $S^{n-1} \hookrightarrow \mathbb{R}^n$, so h is in particular an antipodal map $S^n \to \mathbb{R}^n$.

Similarly, (3) ⇒ (2) follows by defining $h(x) = g(x)/\|g(x)\|$, assuming g is nowhere vanishing.

The version of (4) where all the sets are closed follows from (1) by defining

$$f(x) = (d(x, A_1), \ldots, d(x, A_n))$$

where d denotes Euclidean distance. Then this is a continuous map $S^n \to \mathbb{R}^n$, so find some x with $f(x) = f(-x)$. If one coordinate of $f(x)$ is zero, then we're done. If not, then $x, -x \notin A_1, \ldots, A_n$, so $x, -x \in A_{n+1}$, and we're again done. For the case where some of the sets are open, we simply "fatten" them by ε to closed sets, and then let $\varepsilon \to 0$.

Finally, for (4) ⇒ (3), cover S^{n-1} by $n+1$ sets $B_1, \ldots, B_{n+1}$ that contain no antipodal points (e.g. project the facets of the standard simplex onto S^{n-1}). Then if an antipodal $h : S^n \to S^{n-1}$ existed, then $h^{-1}(B_1), \ldots, h^{-1}(B_{n+1})$ would contradict (4). ☐

Remark 1. Why might we expect the Borsuk-Ulam theorem to be a useful tool? There are many ways to answer this question, but I like to think of it as the "correct" higher-dimensional generalization of the Intermediate Value Theorem, which we hopefully already believe is a useful tool. Indeed, the $n = 1$ case of Borsuk-Ulam is precisely equivalent to the Intermediate Value Theorem; for given an antipodal map $g : S^1 \to \mathbb{R}$, pick some $x \in S^1$. If $g(x) = 0$, we're done. If not, then rotate x by 180°; this is a continuous map $[0, \pi] \to \mathbb{R}$, and at the end of it we've moved form $g(x)$ to $-g(x)$. By the Intermediate Value Theorem, we must have hit zero at some point.

2 First Application: The Kneser Conjecture

Definition 1. Given $n \geq 1$ and $k \leq n/2$, the Kneser graph $KG(n,k)$ has vertex set $V = \binom{[n]}{k}$ and $S, T \in V$ are adjacent if and only if $S \cap T = \emptyset$.

Kneser graphs are very useful (counter)examples in a lot of instances. They have many nice properties, including many symmetries (the automorphism group of $KG(n,k)$ is S_n, since we may permute $[n]$ without changing the graph) and no short odd cycles. Also, you might believe they're important because $KG(5,2)$ is the Petersen graph.

Proposition 1. $\chi(KG(n,k)) \leq n - 2k + 2$.

Proof. Every set that contains the element 1 we color with color 1; they form an independent set since they're never disjoint. Every remaining set that contains 2 is colored with color 2; it forms another independent set. We keep doing this up to element $n - 2k + 1$. Finally, every remaining uncolored set is a subset of $\{n - 2k + 2, \ldots, n\}$, which has size $2k - 1$, so any two k-subsets must intersect. So we can color all the remaining ones with this final color. □

Kneser conjectured that this bound was tight, and this turned out to be quite hard to prove. One reason for the difficulty is that Kneser graphs have a very low *fractional* chromatic number (namely n/k), and many of our techniques for lower-bounding the chromatic number actually lower-bound χ_f.

The Kneser Conjecture was eventually proved by Lovász (1978), in probably the first real application of the Borsuk-Ulam Theorem to combinatorics. There have since been many versions of the proof; the following, due to Greene, is the simplest I know.

Proof. Set $d = n - 2k + 1$, and for $x \in S^d$, let $H(x)$ denote the open hemisphere centered at x. Let $X \subset S^d$ be an n-point set in general position (namely, no $d + 1$ of them lie on a hyperplane through the center of S^d). Fix a bijection $X \leftrightarrow [n]$, so that we think of $V(KG(n,k)) = \binom{X}{k}$. Suppose for contradiction that we could color $KG(n,k)$ with d colors. For $i \in [d]$, let

$$A_i = \left\{ x \in S^d : H(x) \supseteq S \in \binom{X}{k}, \chi(S) = i \right\}$$

Finally, let $A_{d+1} = S^d \setminus (A_1 \cup \cdots \cup A_d)$. Then each A_i is open for $i \in [d]$, while A_{d+1} is closed. By version (4) of the Borsuk-Ulam theorem, there is some $i \in [d+1]$ and $x \in S^d$ such that $x, -x \in A_i$.

First, observe that $i \neq d + 1$. For that would mean that $H(x), H(-x)$ both do not contain a k-set of *any* color, so in both $H(x)$ and $H(-x)$ contain at most $k-1$ points of X. This leaves at least $n-2(k-1) = d+1$ points on the equator separating $H(x), H(-x)$, which contradicts the general position assumption.

So $i \in [d]$. But that means that there are two sets $S_1, S_2 \in \binom{X}{k}$ with $S_1 \subseteq H(x), S_2 \subseteq H(-x)$ and $\chi(S_1) = \chi(S_2) = i$. But S_1 and S_2 must be disjoint, since they lie in separate hemispheres, so they can't receive the same color; a contradiction. □

Remark 2. Matoušek (2004) eventually found a "purely combinatorial" proof of the Kneser Conjecture, but in a fairly uninteresting way: he took the topological proofs and carefully excised all the topology, while maintaining all the ideas from the original proofs.

Remark 3. There have been several generalizations of this result, all of which use a similar topological proof. For instance, Dol'nikov proved a generalization for Kneser graphs of arbitrary set systems (rather than just $\binom{[n]}{k}$), and Schrijver found a vertex-critical subgraph of $KG(n,k)$ that has the same chromatic number.

3 Index and Coindex of $\mathbb{Z}_2$-spaces

Let's begin with some graph-theoretic motivation. For two graphs G, H, we'll write $G \to H$ if there exists some graph homomorphism $f : G \to H$, and $G \nrightarrow H$ if no such homomorphism exists. Then observe that if G is a loopless graph, then

$$\chi(G) = \min\{n : G \to K_n\} \qquad \omega(G) = \max\{n : K_n \to G\}$$

Indeed, a coloring of G with n colors is precisely a homomorphism $G \to K_n$ by mapping each vertex to its color; since every edge goes between distinct color classes, it'll get mapped to an edge of K_n. Similarly, if G is loopless, then every homomorphism $K_n \to G$ must be injective, lest we collapse an edge of K_n to a (non-existent) loop in G. Therefore, a homomorphism $K_n \to G$ precisely identifies a clique subgraph of G, and $\omega(G)$ is the maximal size of such a subgraph. Also note that if we drop the looplessness assumption, then both χ and ω will be infinite in this formulation (since a vertex with a loop can never map to any loopless graph, and we can map all of K_n to such a vertex).

Recall the simple fact that for any graph G, $\omega(G) \le \chi(G)$. In this language, this is equivalent to the statement that $K_n \nrightarrow K_{n-1}$ for any n; for if $\omega(G) > \chi(G)$, then we'd have some n with $K_n \to G \to K_{n-1}$. In fact, this statement also means that the scale according to which we're measuring our graphs, namely the sequence $(K_1, K_2, K_3, \dots)$ is really a well-defined measurement scale, and thus that we get informative quantities.

Once presented in this way, the fact that there is no antipodal map $S^n \to S^{n-1}$ looks very suggestive; can we come up with a topological analogue of the theory of clique and chromatic numbers with the spheres as our measurement scale? To start with, we need to restrict our notion of space and homomorphism so that we can actually apply this fact, since there are many continuous maps $S^n \to S^{n-1}$.

Definition 2. A $\mathbb{Z}_2$-space is a pair (X, ρ), where X is a topological space and $\rho : X \to X$ is a continuous involution (namely a continuous map with $\rho \circ \rho = \mathrm{id}$); note that this is precisely an action of the group $\mathbb{Z}_2$ on X. In case ρ has no fixed points, then we say that this a *free* $\mathbb{Z}_2$-space (as indeed in this case this is a free action).

Definition 3. A $\mathbb{Z}_2$-map (aka antipodal map) between $\mathbb{Z}_2$-spaces $(X, \rho), (Y, \sigma)$ is a continuous map $f : X \to Y$ that is $\mathbb{Z}_2$-equivariant, namely $\sigma \circ f = f \circ \rho$. If there is a $\mathbb{Z}_2$-map from (X, ρ) to (Y, σ), we write $(X, \rho) \xrightarrow{\mathbb{Z}_2} (Y, \sigma)$, and $(X, \rho) \overset{\mathbb{Z}_2}{\nrightarrow} (Y, \sigma)$ if there is no such map.

Definition 4. The index and coindex of a (free) $\mathbb{Z}_2$ space (X, ρ) are defined as

$$\mathrm{ind}_{\mathbb{Z}_2}(X, \rho) = \min\{n : (X, \rho) \xrightarrow{\mathbb{Z}_2} (S^n, -)\}$$
$$\mathrm{coind}_{\mathbb{Z}_2}(X, \rho) = \max\{n : (S^n, -) \xrightarrow{\mathbb{Z}_2} (X, \rho)\}$$

where $(S^n, -)$ denotes the $\mathbb{Z}_2$-space whose underlying space is S^n and whose $\mathbb{Z}_2$-action is given by the antipodal mapping.

Remark 4. Observe that the collection of all $\mathbb{Z}_2$-spaces and $\mathbb{Z}_2$-maps forms a category. The canonical example to keep in mind of a (free) $\mathbb{Z}_2$-space is S^n, equipped with the antipodal mapping. Finally, observe that just as in the case of non-loopless graphs, non-free $\mathbb{Z}_2$-spaces are completely uninteresting from the perspective of index and coindex: we can never map a non-free space to a free one, so the index will be infinite, while we can map any $\mathbb{Z}_2$ space to a fixed point of a non-free space, so its coindex will be infinite as well.

From now on, I will usually not explicitly write ρ when describing a $\mathbb{Z}_2$-space, though it's important to remember that the same topological space may have many different $\mathbb{Z}_2$-structures.

Here are some basic properties of the index and coindex.

Proposition 2.

1. *If* $X \xrightarrow{\mathbb{Z}_2} Y$*, then* $\operatorname{ind}_{\mathbb{Z}_2} X \le \operatorname{ind}_{\mathbb{Z}_2} Y$ *and* $\operatorname{coind}_{\mathbb{Z}_2} X \le \operatorname{coind}_{\mathbb{Z}_2} Y$*.*

2. $\operatorname{ind}_{\mathbb{Z}_2} S^n = \operatorname{coind}_{\mathbb{Z}_2} S^n = n$*.*

3. *If* X *is* $(n-1)$*-connected (namely* $\pi_k(X) = 0$ *for all* $k < n$*), then* $\operatorname{coind}_{\mathbb{Z}_2} X \ge n$*.*

4. *If* X *is a free simplicial* $\mathbb{Z}_2$*-complex (or CW complex) of dimension* n*, then* $\operatorname{ind}_{\mathbb{Z}_2} X \le n$*.*

Sketch of Proof. We won't really need most of these results, so I'll only give an idea of the proofs.

(1) follows directly from the definition: if $n = \operatorname{ind}_{\mathbb{Z}_2} Y$ then $Y \xrightarrow{\mathbb{Z}_2} S^n$, and since $X \xrightarrow{\mathbb{Z}_2} Y$, by composition $X \xrightarrow{\mathbb{Z}_2} S^n$, so $\operatorname{ind}_{\mathbb{Z}_2} X \le n$; an analogous argument establishes the coindex inequality.

(2) is precisely the Borsuk-Ulam Theorem, that $S^n \overset{\mathbb{Z}_2}{\nrightarrow} S^{n-1}$, as indicated above.

For (3), we can construct a $\mathbb{Z}_2$-map $S^k \to X$ inductively for all $0 \le k \le n$. For the base case, pick an arbitrary $x \in X$, and map $S^0 = \{\pm 1\}$ to X by sending $1 \mapsto x, -1 \mapsto \rho(x)$. Inductively, if we have a map $S^k \to X$, we consider S^k as the equator of S^{k-1}. Since X is $(n-1)$-connected, the map $S^k \to X$ is nullhomotopic, so it can be extended (non-uniquely) to the upper hemisphere of S^{k+1}; by reflection, since we require our map to be a $\mathbb{Z}_2$-map, we are forced in how to extend to the lower hemisphere, and we end up with a $\mathbb{Z}_2$-map $S^{k+1} \to X$.

For (4), we can just repeat the proof of (3) once we observe that we never actually used the structure of S^k, just that it has a decomposition into cells that are matched by the $\mathbb{Z}_2$-action. So since S^n is $(n-1)$-connected, the same argument allows us to construct a map $X \xrightarrow{\mathbb{Z}_2} S^n$, so $\operatorname{ind}_{\mathbb{Z}_2} X \le n$. $\square$

Finally, since it was mentioned in the previous proposition, let's recall the definition of a simplicial complex:

Definition 5. A *simplicial complex* on a set V (called the *vertices*) is a collection K of subsets of V that is downwards closed. An element of K is called a *face* or *simplex*, and its *dimension* is 1 less than its size.

We will often also think of a simplicial complex as a topological space, gotten by placing the points of V in general position in a sufficiently high-dimensional Euclidean space, and then putting in a copy of the standard d-simplex Δ^d between each $(d+1)$-tuple that is a face of K. I will be intentionally glib about which of these two interpretations I mean, because they are fundamentally the same.

For a point $x \in \mathsf{K}$, we denote by the *support* of x the smallest (i.e. lowest-dimensional) face of K that contains x.

4 Second Application: A More General Chromatic Lower Bound

The above definitions of index and coindex allow us to generalize the proof of the Kneser conjecture to give a lower bound for the chromatic number of any graph. Note that this bound is frequently not tight or difficult to compute, so it is perhaps not extremely useful in practice. However, I find it illustrative both in that it makes the Kneser proof seem less ad hoc and because it is an instance where the topology is very obviously necessary; one needs the topology to even state the result.

First, we will need some definitions. Recall that $N(v)$ denotes the neighborhood of a vertex v in a graph, and we use the convention that $v \notin N(v)$. For a set $A \subseteq V(G)$, $N(A)$ denotes the set of common neighbors to A, namely $\bigcap_{v \in A} N(v)$; in particular, $N(A) \subseteq V(G) \setminus A$. Finally, if A_1, A_2 are disjoint sets of vertices of G, then let $G[A_1, A_2]$ denote the induced bipartite subgraph of G on A_1, A_2 (namely we discard all vertices not in $A_1 \cup A_2$, and discard all edges internal to A_1 or A_2).

Definition 6. Given a graph G, its *box complex* $\mathsf{B}(G)$ is a simplicial complex whose vertices are $V(G) \times \{1, 2\}$ and whose simplices are

$$\{A_1 \uplus A_2 : A_1, A_2 \subseteq V(G); A_1 \cap A_2 = \emptyset; G[A_1, A_2] \text{ is complete}; N(A_1), N(A_2) \neq \emptyset\}$$

where $A_1 \uplus A_2$ denotes $(A_1 \times \{1\}) \cup (A_2 \times \{2\}) \subseteq V(G) \times \{1, 2\}$.

In other words, the simplices of $\mathsf{B}(G)$ correspond to complete bipartite subgraphs of G. The condition $N(A_1), N(A_2) \neq \emptyset$ is unnecessary in the case when both A_1, A_2 are nonempty (for in that case it follows from $G[A_1, A_2]$ being complete), but is important when one of A_1, A_2 is empty (which we allow).

$\mathsf{B}(G)$ becomes a free $\mathbb{Z}_2$-space when we endow it with the involution that swaps the roles of the two sides: $\rho : (v, 1) \leftrightarrow (v, 2)$. Note that this is a free action because we required that A_1, A_2 be disjoint, so $\rho(A_1 \uplus A_2) = A_2 \uplus A_1$ will be disjoint from $A_1 \uplus A_2$.

Example 1. If P_2 denotes the path graph on 3 vertices, then $\mathsf{B}(P_2)$ is a disjoint union of two triangles (2-simplices).

We claim that $\mathsf{B}(K_n)$ is a the cross-polytope of dimension n with two antipodal $(n-1)$-simplices removed. Indeed, recall that the cross-polytope is a simplicial complex with vertex set $\Xi = \{\pm 1, \pm 2, \ldots, \pm n\}$ and whose simplices are all subsets of Ξ that don't contain both i and $-i$ for any $i \in [n]$ (this precisely corresponds to taking the boundary of the convex hull of the vectors $\{\pm e_1, \ldots, \pm e_n\} \subset \mathbb{R}^n$, where $\{e_i\}$ is the standard basis). On the other hand, $\mathsf{B}(K_n)$ has vertex set $[n] \times \{1, 2\}$. Its simplices are all $A_1 \uplus A_2$ with $A_1 \cap A_2 = \emptyset$, *except* for the two $[n] \uplus \emptyset$ and $\emptyset \uplus [n]$, since the set of all vertices of K_n do not have a common neighbor. By using the bijection $\Xi \leftrightarrow [n] \times \{1, 2\}$ that identifies $i \leftrightarrow (i, 1), -i \leftrightarrow (i, 2)$, we see that $\mathsf{B}(K_m)$ is indeed just the cross-polytope with two antipodal faces removed.

Since the cross-polytope is homeomorphic to S^{n-1}, removing these two faces gives us a deformation retraction $\mathsf{B}(K_n) \to S^{n-2}$; moreover, this map is a $\mathbb{Z}_2$-map, since we can deform antipodal points consistently. Similarly, we can embed S^{n-2} as the equator of $\mathsf{B}(K_n)$, so we find that

$$n - 2 \leq \operatorname{coind}_{\mathbb{Z}_2} \mathsf{B}(K_n) \leq \operatorname{ind}_{\mathbb{Z}_2} \mathsf{B}(K_n) \leq n - 2$$

and thus they both equal $n - 2$.

One of the most important properties of the box complex construction is that it is *functorial*:

Definition 7. Let $f : G \to H$ be a graph homomorphism. We construct a simplicial map $\mathsf{B}(f) : \mathsf{B}(G) \to \mathsf{B}(H)$ by mapping

$$(v, i) \mapsto (f(v), i)$$

where $v \in V(G), i \in \{1, 2\}$. Then it is a simplicial map because a complete bipartite subgraph of G will be mapped to a complete bipartite subgraph of H under f; moreover, it is a $\mathbb{Z}_2$-map because the $\mathbb{Z}_2$-action in all instances is just to swap 1 and 2 in the second coordinate. Thus, B gives us a functor from the category of graphs to the category of (free) $\mathbb{Z}_2$-spaces.

Corollary 1. *For any graph G,*

$$\chi(G) \geq \operatorname{ind}_{\mathbb{Z}_2} \mathsf{B}(G) + 2$$

Proof. Let $r = \chi(G)$. Then there is a graph homomorphism $G \to K_r$. Applying B to this gives us $\mathsf{B}(G) \xrightarrow{\mathbb{Z}_2} \mathsf{B}(K_r)$, so

$$\operatorname{ind}_{\mathbb{Z}_2} \mathsf{B}(G) \leq \operatorname{ind}_{\mathbb{Z}_2} B(K_r) = r - 2$$

or

$$\chi(G) = r \geq \operatorname{ind}_{\mathbb{Z}_2} \mathsf{B}(G) + 2$$

□

This is a very nice and completely general lower bound on the chromatic number, though one could argue that the box complex is a sort of artificial construction. A more natural simplicial complex associated to a graph is the following:

Definition 8. Given a graph G, its *neighborhood complex* $\mathsf{N}(G)$ is the simplicial complex whose vertices are $V(G)$ and $A \subseteq V(G)$ is a simplex if and only if $N(A) \neq \emptyset$; in other words, all neighborhoods of vertices are top-dimensional faces of $\mathsf{N}(G)$.

Theorem 2 (Lovász, 1978)**.** *If $\mathsf{N}(G)$ is k-connected, then*

$$\chi(G) \geq k+3$$

Idea of proof. What we want to say is that since $\mathsf{N}(G)$ is k-connected, Proposition 2 implies that $\mathrm{ind}_{\mathbb{Z}_2} \mathsf{N}(G) \geq k+1$. Then if we can construct a map $\mathsf{N}(G) \xrightarrow{\mathbb{Z}_2} \mathsf{B}(G)$, then we will conclude that $\mathrm{ind}_{\mathbb{Z}_2} \mathsf{B}(G) \geq k+1$ as well, so by the above corollary $\chi(G) \geq k+3$.

This turns out to basically work, except for an obvious flaw: $\mathsf{N}(G)$ is not endowed with a $\mathbb{Z}_2$-action, so nothing in the previous paragraph makes any sense. The way to solve this is to construct yet another simplicial complex $\mathsf{L}(G)$ that is a deformation retraction of $\mathsf{N}(G)$ and that is a $\mathbb{Z}_2$-space, and then to proceed as above (using the fact that since $\mathsf{N}(G)$ is k-connected and a deformation retraction is a homotopy equivalence, $\mathsf{L}(G)$ will be k-connected as well). I'll skip the details, since they're fairly technical and not particularly enlightening. □

Remark 5. This theorem is how Lovász originally proved the Kneser Conjecture; he first proved the above directly from the Borsuk-Ulam Theorem, and then proved that the neighborhood complex of $KG(n,k)$ was highly connected, something that is actually not so simple to show.

However, as promised, this is a generalization of the proof of the Kneser Conjecture that provides a general chromatic lower bound, and whose statement fundamentally uses topological notions. Also notice that the key part of the proof we did was actually extremely straightforward: once B is a functor, then the fact that the index and chromatic number are defined analogously immediately implies we can obtain results if we know $\mathrm{ind}_{\mathbb{Z}_2} \mathsf{B}(K_n)$.

5 Joins, Deleted Joins, and Embeddings

So far, we have proved existence and non-existence of various maps, and have used this to get information on the $\mathbb{Z}_2$-index and coindex of various spaces. Now, we will try to reverse this: by bounding the index and index of some spaces, we will conclude that no $\mathbb{Z}_2$-map between them exists.

We begin with some more topological notions:

Definition 9. Given two topological spaces X, Y, their *join* $X * Y$ is the space

$$X * Y = (X \times Y \times [0,1]) / \sim$$

where $\sim$ is the equivalence relation that identifies all points of x at time 0, and all points of y at time 1:

$$(x, y, 0) \sim (x', y, 0) \qquad (x, y, 1) \sim (x, y', 1)$$

for all $x, x' \in X, y, y' \in Y$. We write a point in $X * Y$ as a "formal convex combination" $tx \oplus (1-t)y$ for $x \in X, y \in Y$; this notation stresses that at $t = 0$, what x we pick doesn't matter, and at $t = 1$, what y we pick doesn't matter.

This notation also reflects a deeper truth: if X, Y are bounded subsets of $\mathbb{R}^n$ that lie on *skew* hyperplanes (i.e. non-parallel, non-intersecting hyperplanes), then $X * Y$ is the subset of $\mathbb{R}^n$ consisting of all line segments with one endpoint in x and the other in y, namely

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

If K, L are simplicial complexes, then we can also define their join $\mathsf{K} * \mathsf{L}$ as the simplicial complex whose vertex set is $V(\mathsf{K}) \uplus V(\mathsf{L})$ and whose simplices are

$$\{F \uplus G : F \in \mathsf{K}, G \in \mathsf{L}\}$$

Using the above geometric construction, it's not too hard to see that these two notions coincide.

Example 2. If I is the interval $[0, 1]$, then its join with itself is a 3-simplex Δ^3:

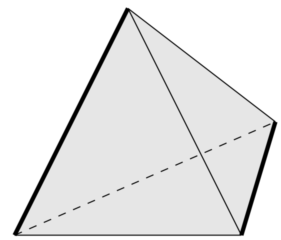

One way to generalize this is to observe that if Δ^0 denotes a one-point space, then $(\Delta^0)^{*n}$ is the simplex Δ^{n-1}; this follows from the geometric interpretation, which tells us that $(\Delta^0)^{*n}$ should be the convex hull of n points in a sufficiently high-dimensional space, which is precisely Δ^{n-1}. From this, we can inductively obtain that $\Delta^m * \Delta^n = \Delta^{m+n+1}$.

Recall that S^0 is a two-point discrete space. Then for any X, $X * S^0$ is the *suspension* of X, which consists of a double cone over X. In particular, $S^0 * S^n = S^{n+1}$, which implies by induction that $S^m * S^n = S^{m+n+1}$.

Given any space X, its two-fold join $X^{*2} = X * X$ naturally comes with a $\mathbb{Z}_2$-action that simply swaps the two sides:

$$tx \oplus (1-t)y \mapsto (1-t)y \oplus tx$$

Note that this is very similar to the $\mathbb{Z}_2$ action on the box complex we saw earlier. However, this is not a free $\mathbb{Z}_2$-action, since any point $\frac{1}{2}x \oplus \frac{1}{2}x$ is a fixed point. This motivates the following definition:

Definition 10. For any topological space X, its *deleted join* is the space

$$X^{*2}_\delta = (X * X) \setminus \{\tfrac{1}{2}x \oplus \tfrac{1}{2}x : x \in X\}$$

The δ denotes that we've deleted the "diagonal" $\{\frac{1}{2}x \oplus \frac{1}{2}x : x \in X\}$. This is a free $\mathbb{Z}_2$-space with the action given by exchanging the two sides, as above.

For simplicial complexes, we have a similar but not entirely identical notion:

Definition 11. For any simplicial complex K, its *deleted join* is the simplicial complex K^{*2}_Δ whose vertices are $V(\mathsf{K}) \times \{1, 2\}$ and whose simplices are

$$\{F_1 \uplus F_2 : F_1, F_2 \in \mathsf{K}, F_1 \cap F_2 = \emptyset\}$$

It is a subcomplex of K^{*2}. Note that we delete strictly more than above, since we delete the whole offending simplex. Also note that K^{*2}_Δ is a free simplicial $\mathbb{Z}_2$-space, with $\mathbb{Z}_2$-action given by swapping the two sides; just as in the case of the box complex, this is free precisely because we delete all simplices whose two sides intersect.

Example 3. Let I denote the interval $[0, 1]$. When thought of as a topological space, its two-fold join I^{*2} is a 3-simplex Δ^3. Its diagonal $\{\frac{1}{2}x \oplus \frac{1}{2}x : x \in I\}$ is the red line in the following picture:

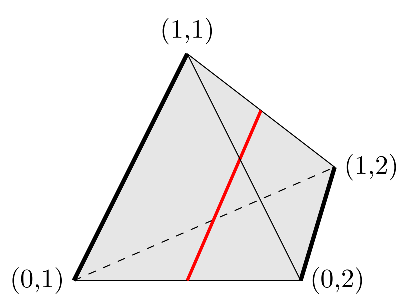

Thus, we see that I^{*2} is, topologically, a 3-ball with a diameter removed. This deformation retracts to S^1, and this deformation retraction is compatible with the $\mathbb{Z}_2$ structures.

If we think of I as a simplicial complex with two vertices $0, 1$ and one 1-simplex $\{0, 1\}$, then I^{*2}_Δ is a simplicial complex with four vertices $\{(0, 1), (1, 1), (0, 2), (1, 2)\}$. It has only four simplices, as depicted in the following picture (the simplices that have been deleted are dotted):

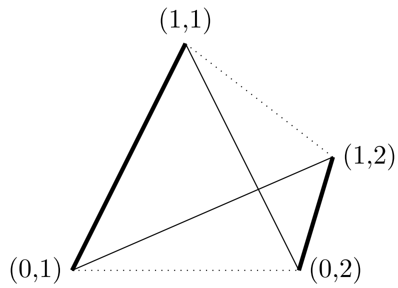

Thus, we see that $I^{*2}_\Delta \cong S^1$.

We end with two simple lemmas about properties of these deleted joins.

Lemma 1. *If* K, L *are simplicial complexes, then*

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

In other words, the deleted join commutes with the ordinary join.

Proof. Both sides have vertex set $(V(\mathsf{K}) \uplus V(\mathsf{L})) \times \{1, 2\}$. A simplex on the left-hand side is of the form $(F_1 \uplus G_1) \uplus (F_2 \uplus G_2)$ for $F_1, F_2 \in \mathsf{K}, G_1, G_2 \in \mathsf{L}$, with the condition that $(F_1 \uplus G_1) \cap (F_2 \uplus G_2) = \emptyset$. This condition is equivalent to $F_1 \cap F_2 = G_1 \cap G_2 = \emptyset$. But that is precisely the condition for $(F_1 \uplus G_1) \uplus (F_2 \uplus G_2)$ being a simplex of $\mathsf{K}^{*2}_\Delta * \mathsf{L}^{*2}_\Delta$. □

Lemma 2. $\operatorname{ind}_{\mathbb{Z}_2}(\mathbb{R}^n)^{*2}_\delta \leq n$.

Proof. We can exhibit a $\mathbb{Z}_2$-map $(\mathbb{R}^n)^{*2}_\delta \to S^n$. First, let $f : (\mathbb{R}^n)^{*2}_\delta \to \mathbb{R}^{2n+2}$ be given by

$$f(t\vec{x} \oplus (1-t)\vec{y}) = (t, \vec{x}, 1-t, \vec{y})$$

for $\vec{x}, \vec{y} \in \mathbb{R}^n, t \in [0, 1]$. We can view $\mathbb{R}^{2n+2}$ as $(\mathbb{R}^{n+1})^2$; then f never hits the diagonal $D = \{(\vec{z}, \vec{z}) : \vec{z} \in \mathbb{R}^{n+1}\}$. For if $f(t\vec{x} \oplus (1-t)\vec{y}$ hit the diagonal, then we would necessarily have $t = 1-t$ and $\vec{x} = \vec{y}$, and these are precisely the points we removed from $(\mathbb{R}^n)^{*2}_\delta$. Now, $(\mathbb{R}^{n+1})^2 \setminus D$ naturally has a free $\mathbb{Z}_2$-action given by exchanging the two coordinates, and this turns f into a $\mathbb{Z}_2$-map.

Finally, we have a $\mathbb{Z}_2$-map $g : (\mathbb{R}^{n+1})^2 \setminus D \to S^n$ defined by

$$g(\vec{w}, \vec{z}) = \frac{\vec{w} - \vec{z}}{\|\vec{w} - \vec{z}\|} \in S^n \subseteq \mathbb{R}^{n+1}$$

This is well-defined since we deleted the diagonal, and is a $\mathbb{Z}_2$-map by construction. Composing $g \circ f$ shows that $(\mathbb{R}^n)^{*2}_\delta \xrightarrow{\mathbb{Z}_2} S^n$, as claimed[1].

Note that this is more or less the same as the $\mathbb{Z}_2$-map $I^{*2}_\delta \to S^1$ that we explicitly drew earlier. □

[1] In fact, this proof is wrong, even though the statement is correct. The error comes from the fact that f is not well-defined when $t \in \{0, 1\}$. In order to fix this, one has to first embed two copies of $\mathbb{R}^n$ as balls in skew hyperplanes in $\mathbb{R}^{2n+2}$ and then use the honest convex combinations instead of the formal ones.

6 Third Application: Non-Planarity and the Topological Radon Theorem

As an immediate application of these tools, let's prove that $K_{3,3}$ is non-planar. For this, suppose that we had a map $f : K_{3,3} \to \mathbb{R}^2$ with the property that the images of distinct edges never intersect. Then we can define a new map $f^{*2} : (K_{3,3})^{*2}_{\Delta} \to (\mathbb{R}^2)^{*2}_{\delta}$ by

$$f^{*2}(tx \oplus (1-t)y) = tf(x) \oplus (1-t)f(y)$$

To see that this indeed lands in $(\mathbb{R}^2)^{*2}_{\delta}$, suppose not. Then we must hit a point of the form $\frac{1}{2}z \oplus \frac{1}{2}z$ for $z \in \mathbb{R}^2$, which implies that there are x, y so that $f(x) = f(y) = z$. x and y cannot lie in the same simplex, for we deleted such points from $(K_{3,3})^{*2}_{\Delta}$, so they must have disjoint supports. But then that contradicts the assumption that f never allows the images of distinct edges to intersect. Finally, note that f^{*2} is a $\mathbb{Z}_2$-map, so we've found that $(K_{3,3})^{*2}_{\Delta} \xrightarrow{\mathbb{Z}_2} (\mathbb{R}^2)^{*2}_{\delta}$.

Observe that $K_{3,3} = \mathsf{D}_3 * \mathsf{D}_3$, where D_3 is a discrete 3-point simplicial complex (i.e. with no simplices of positive dimension). Therefore, by Lemma 1,

$$(K_{3,3})^{*2}_{\Delta} = (\mathsf{D}_3 * \mathsf{D}_3)^{*2}_{\Delta} = (\mathsf{D}_3)^{*2}_{\Delta} * (\mathsf{D}_3)^{*2}_{\Delta}$$

Here is $(\mathsf{D}_3)^{*2}_{\Delta}$, which is just $K_{3,3}$ with a matching deleted.

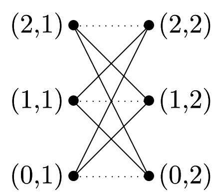

By unraveling it, we can see that it's topologically S^1. Since $S^1 * S^1 = S^3$, we see in particular that $\operatorname{ind}_{\mathbb{Z}_2}(K_{3,3})^{*2}_{\Delta} = 3$. On the other hand, we saw above that $\operatorname{ind}_{\mathbb{Z}_2}(\mathbb{R}^2)^{*2}_{\delta} \leq 2$, so $(K_{3,3})^{*2}_{\Delta} \not\xrightarrow{\mathbb{Z}_2} (\mathbb{R}^2)^{*2}_{\delta}$. This contradicts the existence of f^{*2} above, so $K_{3,3}$ is non-planar.

We can generalize this example as follows:

Theorem 3. *For any simplicial complex* K, *if*

$$\operatorname{ind}_{\mathbb{Z}_2} \mathsf{K}^{*2}_{\Delta} > n$$

then K *cannot be embedded in* $\mathbb{R}^n$. *In fact, any continuous map* $f : \mathsf{K} \to \mathbb{R}^n$ *will necessarily identify two points with disjoint supports in* K.

Proof. The proof is identical to what we did above. If there were a map $f : \mathsf{K} \to \mathbb{R}^n$ without this property, then we could construct a $\mathbb{Z}_2$-map $f^{*2} : \mathsf{K}^{*2}_{\Delta} \to (\mathbb{R}^n)^{*2}_{\delta}$ in precisely the same way as above. However, by assumption and by our earlier calculations, $\operatorname{ind}_{\mathbb{Z}_2} \mathsf{K}^{*2}_{\Delta} > \operatorname{ind}_{\mathbb{Z}_2}(\mathbb{R}^n)^{*2}_{\delta}$, so $\mathsf{K}^{*2}_{\Delta} \not\xrightarrow{\mathbb{Z}_2} (\mathbb{R}^n)^{*2}_{\delta}$; a contradiction.

The basic idea in this proof is that we have one tool for proving that a certain sort of map doesn't exist, namely the $\mathbb{Z}_2$-index. We want to show that another sort of map (an embedding) doesn't exist. So we build new spaces whose structure is such that any $\mathbb{Z}_2$-map between them corresponds to an embedding between the original spaces, and then use our one tool to show that can't happen. $\square$

To see a more exciting application of this result than the non-planarity of $K_{3,3}$, recall the classical Radon Theorem from convex geometry:

Theorem 4 (Radon)**.** *Given any* $n+2$ *points in* $\mathbb{R}^n$, *they can be partitioned into two sets whose convex hulls intersect.*

Proof. We can prove this by induction on n (this is not the standard proof one finds in most books, but it is one that I find enlightening). The $n = 1$ case is simple: given 3 points in $\mathbb{R}^1$, they come in some order. Then the middle point is in the convex hull of the other two.

For the induction step, say we have $n + 2$ points in $\mathbb{R}^n$. One can show that either all of them lie on a hyperplane (in which case we're done by induction), or else there is some hyperplane H containing n of them that separates the other two, which we call a, b. Then the convex hull of a and b (a line segment) intersects H at some point. Let this intersection point be c. By induction, we can partition c and the other n points on H into two sets whose convex hulls intersect; remove c from whichever of these two sets it lies in and replace it by a and b. Then the convex hulls still intersect, since c is in the convex hull of a, b. $\square$

To make the connection with embeddings more obvious, here is an equivalent statement of Radon's Theorem:

Theorem 5 (Radon, Take 2)**.** *For any affine*[2] *map $\Delta^{n+1} \to \mathbb{R}^n$, there are two disjoint faces whose images intersect.*

This is equivalent to the previous statement: an affine mapping is precisely determined by where it places the $n + 2$ vertices in $\mathbb{R}^n$, and any convex hull of a subset is the affine image of a face of Δ^{n+1}.

Using our techniques, we can prove a major generalization.

Theorem 6 (Topological Radon Theorem, Bajmóczy-Bárány 1979)**.** *For any* continuous *map $f : \Delta^{n+1} \to \mathbb{R}^n$, there are two disjoint faces whose images intersect.*

Proof. We apply Theorem 3. To do so, we need to show that $\text{ind}_{\mathbb{Z}_2}(\Delta^{n+1})^{*2}_{\Delta} > n$. Recall that $\Delta^{n+1} = (\Delta^0)^{*(n+2)}$. Therefore, by Lemma 1,

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

and in particular $\text{ind}_{\mathbb{Z}_2}(\Delta^{n+1})^{*2}_{\Delta} = n + 1 > n$, as desired. $\square$

Remark 6. There is a "density" version of the Radon theorem, originally due to Boros and Füredi, which was also made topological by Gromov and others; such results have been extremely influential in the recent study of high-dimensional expanders.

Finally, let me mention a beautiful result that connects both embeddability results and chromatic number results:

Theorem 7 (Sarkaria 1991)**.** *Let K be any simplicial complex on n vertices, and let $\mathcal{F}$ denote the set system of all inclusion-minimal non-faces of K. Then*

$$\text{ind}_{\mathbb{Z}_2}(\mathsf{K}^{*2}_{\Delta}) + \chi(KG(\mathcal{F})) \geq n - 1$$

where $KG(\mathcal{F})$ is the Kneser graph generated by $\mathcal{F}$. Note that this can be used both for embeddability lower bounds and chromatic number lower bounds!

7 G-spaces, G-maps, and G-indices

As we observed, what we call a $\mathbb{Z}_2$-space is simply a topological space together with a (continuous) action of the group $\mathbb{Z}_2$. Seeing as this was very useful, it makes sense to extend the definition to arbitrary groups; for concreteness, we stick with finite groups.

[2] An affine map is one that picks images for the vertices and then extends affine-linearly over all the faces.

Definition 12. For a finite group G, a G-space is a tuple (X, Φ) where X is a topological space and Φ is a map $G \to \operatorname{Aut} X$, where $\operatorname{Aut} X$ denotes the space of all homeomorphisms $X \to X$. Equivalently, for each $g \in G$ we have a continuous map $\varphi_g : X \to X$ which commutes with the group operation, namely $\varphi_g \circ \varphi_h = \varphi_{gh}$.

In the case where $G = \mathbb{Z}_n$ is cyclic, then a G-space is simply (X, ν), where $\nu : X \to X$ is a continuous map satisfying $\nu^n = \mathrm{id}$.

A G-space is called *free* if for any $1 \neq g \in G$, φ_g has no fixed points.

Definition 13. If $(X, \Phi), (Y, \Psi)$ are G-spaces, then a G-map is a continuous equivariant map $f : X \to Y$, namely a continuous map with $\psi_g \circ f = f \circ \varphi_g$ for all $g \in G$. As before, we write $X \xrightarrow{G} Y$ if a G-map $X \to Y$ exists, and $X \overset{G}{\not\to} Y$ otherwise.

In order to mimic the $\mathbb{Z}_2$-index we previously developed, we need to come up with a suitable measurement scale; a sequence $X_1, X_2, \ldots$ of G-spaces with the property that $X_n \overset{G}{\not\to} X_{n-1}$ for all n. Unlike previously, where we could just take the spheres, there is no immediately obvious candidate. One concrete choice (out of several possible, but all equivalent) is the following:

Definition 14. Observe that G, endowed with the discrete topology, is a free G-space, with G acting by left multiplication. This induces (though this requires some check) a free G-action on the n-fold join G^{*n}. Thus, we define the space $E_n G = G^{*(n+1)}$, and this will serve as our measuring scale.

Observe that in the case $G = \mathbb{Z}_2$, we actually recover the spheres, since $\mathbb{Z}_2 = S^0$. This also explains the slightly strange indexing $E_n G = G^{*(n+1)}$.

Definition 15. The G-index and coindex are defined as

$$\operatorname{ind_G} X = \min\{n : X \xrightarrow{G} E_n G\} \qquad \operatorname{coind_G} X = \max\{n : E_n G \xrightarrow{G} X\}$$

It turns out to be surprisingly difficult to show that $E_n G \overset{G}{\not\to} E_{n-1} G$, though it is true. In particular, it is harder than proving the Borsuk-Ulam theorem, which is just a special case of it.

The theory of $\mathbb{Z}_p$-index allows us to prove a nice generalization of the Borsuk-Ulam Theorem:

Theorem 8. *Let p be prime. Let (X, ν) be a $\mathbb{Z}_p$-space with $\operatorname{ind}_{\mathbb{Z}_p} X \geq n(p-1)$. Then for any continuous map $f : X \to \mathbb{R}^n$, there is some $x \in X$ with*

$$f(x) = f(\nu(x)) = \cdots = f(\nu^{p-1}(x))$$

In other words all orbits of x are mapped to the same point.

Proof. The proof, like many of the ones we saw before, goes by arguing that if such a map did not exist, then we could get a $\mathbb{Z}_p$-map from a space of some $\mathbb{Z}_p$-index to a space of lower $\mathbb{Z}_p$-index.

Specifically, suppose we had an f for which this were not true. From f we can construct a map $f^\times : X \to \mathbb{R}^{np}$ defined by

$$f^\times(x) = (f(x), f(\nu(x)), \ldots, f(\nu^{p-1})(x))$$

Observe that $f^\times$ is a $\mathbb{Z}_p$-map, if we endow $\mathbb{R}^{np}$ the $\mathbb{Z}_p$-action given by cyclically permuting the coordinates. By the assumption on f, we have that the codomain of $f^\times$ is actually

$$Y = \mathbb{R}^{np} \setminus \{(z, z, \ldots, z) : z \in \mathbb{R}^n\}$$

On the other hand, it is a simple computation that

$$\operatorname{ind}_{\mathbb{Z}_p} Y \leq n(p-1) - 1$$

which contradicts the assumption that $\operatorname{ind}_{\mathbb{Z}_p} X \geq n(p-1)$. $\square$

One can also generalize the theory of deleted joins; here are two basic definitions:

Definition 16. If X is a topological space, its *n-fold deleted join* is

$$X_{\delta}^{*n} = X^{*n} \setminus \{\tfrac{1}{n}x \oplus \cdots \oplus \tfrac{1}{n}x : x \in X\}$$

If K is a simplicial complex, its *pairwise n-fold deleted join* is the simplicial complex

$$\mathsf{K}_{\Delta(2)}^{*n} = \{F_1 \uplus \cdots \uplus F_n \in K^{*n} : F_1, \ldots, F_n \text{ are pairwise disjoint}\}$$

These can both be made into $\mathbb{Z}_n$-spaces via cyclically permuting the coordinates. However, if n is not prime, then these will *not* be free $\mathbb{Z}_n$-spaces.

Finally, here is an important lemma:

Lemma 3. *If p is a prime, then*

$$\operatorname{ind}_{\mathbb{Z}_p}((\mathbb{R}^n)_{\delta}^{*p}) \leq (n+1)(p-1)+1$$

8 Fourth Application: The Topological Tverberg Theorem

There is a very important and far-reaching generalization of the Radon Theorem:

Theorem 9 (Tverberg, 1966)**.** *For any $n \geq 1, r \geq 2$, let X be a set of $(n+1)(r-1)+1$ in $\mathbb{R}^n$. Then there is a partition $X = X_1 \sqcup \cdots \sqcup X_r$ so that*

$$\operatorname{conv}(X_1) \cap \operatorname{conv}(X_2) \cap \cdots \cap \operatorname{conv}(X_r) \neq \emptyset$$

Note that the $r = 2$ case is just Radon's Theorem. This theorem itself was greatly generalized by Bárány, Shlosman, and Szűcs (1981) in the case where p is prime:

Theorem 10 (Topological Tverberg)**.** *Let $n \geq 1$ and $p \geq 2$ be a prime. Set $N = (n+1)(p-1)$. Then for any continuous map $f : \Delta^N \to \mathbb{R}^n$, there are p disjoint faces $F_1, \ldots, F_p$ of Δ^N so that*

$$f(F_1) \cap f(F_2) \cap \cdots \cap f(F_r) \neq \emptyset$$

Proof. We mimic the proof of the Topological Radon Theorem. Suppose that there is an f that does not have such an intersection. Then the p-fold join f^{*p} will be a map

$$f^{*p} : (\Delta^N)_{\Delta(2)}^{*p} \xrightarrow{\mathbb{Z}_p} (\mathbb{R}^n)_{\delta}^{*p}$$

As stated above, $\operatorname{ind}_{\mathbb{Z}_p}((\mathbb{R}^n)_{\delta}^{*p}) \leq (n+1)(p-1)+1 = N-1$. So it remains to compute the $\mathbb{Z}_p$-index of $(\Delta^N)_{\Delta(2)}^{*p}$. As in the case of 2-fold deleted joins, the deleted join commutes with the ordinary join: for any simplicial complexes K, L,

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

Therefore, since we can write $\Delta^N = (\Delta^0)^{*(N+1)}$, we find that

$$(\Delta^N)_{\Delta(2)}^{*p} = ((\Delta^0)_{\Delta(2)}^{*p})^{*(N+1)}$$

We can see that $(\Delta^0)_{\Delta(2)}^{*p}$ is a p-point discrete space. In particular, it's equal to $E_0\mathbb{Z}_p$, so its $(N+1)$th join is just $E_N\mathbb{Z}_p$. Thus,

$$\operatorname{ind}_{\mathbb{Z}_p}((\Delta^N)_{\Delta(2)}^{*p}) = N > \operatorname{ind}_{\mathbb{Z}_p}((\mathbb{R}^n)_{\delta}^{*p})$$

which contradicts the existence of f^{*p}. □

Remark 7. The proof of the Topological Tverberg Theorem was eventually extended to the case where r is not just a prime, but a prime power, by Özaydin (1987). So everyone expected the case for general r to eventually be resolved; however, it was shown by Frick in 2014 that it is *false* for general r!