COMBINATORIAL PERSPECTIVES ON BORSUK-ULAM AND BROUWER

MAX WEINSTEIN

Abstract. The Borsuk-Ulam Theorem and Brouwer's Fixed Point Theorem are classic results in topology, with wide-reaching applications. In this paper, we discuss these theorems, and two combinatorial results which are equivalent to these theorems, in the hope of shedding some light on the nature of the classic results.

Contents

1. Description

Here is the structure of the results we will lay out:

$$\begin{tikzcd}
\text{Borsuk-Ulam Theorem} \arrow[r, "\text{p.}11"] \arrow[d, Leftrightarrow, "\text{p.}17"] & \text{Brouwer's Fixed Point Theorem} \arrow[d, Leftrightarrow, "\text{p.}22"] \\
\text{Tucker's Lemma} \arrow[r, "\text{p.}25"] & \text{The Hex Theorem}
\end{tikzcd}$$

Over the course of this paper, we will show each of the arrows on its own, and show each result on its own. The point of these equivalences and implications is not to prove the other results by proving one or two of them. Rather, the point is to give an equivalent reframing of the relationship between two topological results in terms of two combinatorial results.

Readers may already be aware of Sperner's Lemma as a combinatorial statement equivalent to Brouwer's Fixed Point Theorem. The Hex Theorem and Sperner's Lemma are vaguely similar, but the Hex Theorem is a result for which it's more intuitive why one might expect it to relate to Brouwer's Fixed Point Theorem.

We will start with some preliminaries to establish notation and ensure readers are familiar with the topological results used in the paper. This can be skipped if the reader has studied homology theory. We will then demonstrate the top row of the diagram, followed by the equivalence arrows, and finally the bottom row.

2. Preliminaries

Before discussing simplices and homology at length, there is a graph-theoretic result we should get out of the way. This result will be used in the proofs of Tucker's Lemma and the Hex Theorem.

Theorem 2.1. *If every vertex in a graph G has degree at most 2 (that is, at most 2 edges are coming out of each vertex), then G is composed of simple paths, simple cycles, and isolated vertices.*

Proof. We proceed by induction on the number of edges in G.

If G has only one edge, then it either connects two distinct vertices, or connects one vertex to itself. In either case, the edge forms a simple path or cycle and the vertices not on the edge are isolated vertices.

Suppose G has $n+1$ vertices, and the theorem is true for graphs with n vertices. Consider an edge e in G, and let G' be the graph G, with the edge e removed. Removing an edge can only lower the degree of the vertices in a graph, so every vertex in G' has degree at most 2. By our induction hypothesis, G' is composed of simple paths, simple cycles, and isolated vertices. The vertices on e had degree at most 2 in G, so they have degree at most 1 in G'. Each of the vertices on e is thus either an isolated vertex or an end of a simple path in G'. In any of these cases, adding the edge e back into G' produces a graph which is still composed of simple paths, simple cycles, and isolated vertices. But adding e back into G' just gives the graph G. □

Importantly, if such a graph contains a vertex with degree 1, then that vertex is one end of a simple path, and so there must be another vertex with degree 1, at the other end of the path.

2.1. Simplices.

A simplex is the simplest convex shape in a given dimension, like a triangle or a tetrahedron. Every result in this paper uses simplices, so we'll spend some time defining the notation we'll use. The notation comes from [1] and [2].

Definition 2.2. A collection of points $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k \in \mathbb{R}^n$ is **affinely dependent** if there are coefficients α_i, not all zero, such that $\sum_{i=0}^{k} \alpha_i \boldsymbol{v}_i = \boldsymbol{0}$ and $\sum_{i=0}^{k} \alpha_i = 0$. Equivalently, the collection is affinely dependent if the system $\{\boldsymbol{v}_1 - \boldsymbol{v}_0, \ldots, \boldsymbol{v}_n - \boldsymbol{v}_0\}$ is linearly dependent.

A collection of points which is not affinely dependent is called **affinely independent**.

Definition 2.3. The **convex hull** of a collection of points $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k \in \mathbb{R}^n$ is the set

$$\operatorname{conv}\{\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k\} \coloneqq \left\{ \sum_{i=0}^{k} \alpha_i \boldsymbol{v}_i \ : \ \sum_{i=0}^{k} \alpha_i = 1,\ \alpha_i \geq 0 \right\}$$

For every point in the convex hull, the sequence $\{\alpha_i\}$ is called the **coordinates** of the point.

Definition 2.4. A **simplex** is the convex hull of a collection of affinely independent points. The points are called the **vertices** of the simplex, and the **dimension** of the simplex is one less than the number of vertices.

Often, we will just associate a simplex with its vertices, and write it as the set of vertices, with square brackets, like $[\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k]$. We will sometimes need to denote a generic n-simplex (a simplex of dimension n). So we let $\Delta^n \coloneqq [\boldsymbol{0}, \boldsymbol{e}_1, \ldots, \boldsymbol{e}_n]$ be a generic simplex.

Definition 2.5. A **face** of a simplex is the convex hull of any subset of the vertices. The **relative interior** of a simplex is the simplex with all the faces of smaller dimension removed.

The usefulness of simplices is in how they allow us to decompose a space into simple shapes. This is done with a "simplicial complex":

Definition 2.6. A **simplicial complex** is a nonempty family Δ of simplices such that:

(1) Every face of a simplex $\sigma \in \Delta$ is also in Δ, and
(2) the intersection $\sigma_1 \cap \sigma_2$ of any two simplices in Δ is a face of both σ_1 and σ_2.

The **vertex set** $V(\Delta)$ of a simplicial complex is the set of all 0-dimensional simplices in Δ. Equivalently, it's the set of all the vertices of simplices in Δ.

The **polyhedron** of a simplicial complex, denoted $\|\Delta\|$, is the union of all the simplices in Δ.

Because the intersection of two simplices in a complex must be a face of both simplices, every point in $\|\Delta\| \setminus V(\Delta)$ is in the relative interior of exactly one simplex.

Definition 2.7. A **triangulation** of a space X is a simplicial complex Δ such that $\|\Delta\|$ is homeomorphic to X.

So far, we've given a geometric definition of simplices and simplicial complexes. However, there's a way to define them more abstractly. An abstract simplex can generally just be thought of as a set, and a face as a subset. Then a simplicial complex becomes:

Definition 2.8. An **abstract simplicial complex** is a pair (V, K), where V is a vertex set and $K \subseteq 2^V$ such that every subset of an element of K is also in K.

If there is a geometric complex Δ with a bijection $f : \Delta \to K$ which commutes with intersection ($f(\sigma_1 \cap \sigma_2) = f(\sigma_1) \cap f(\sigma_2)$), we call Δ a **geometric realization** of (V, K).

The simplicial complex (V, K) will often be named after the set of simplices, i.e. "K", and the set of vertices called "$V(K)$".

We can always geometrically realize an abstract complex by finding an appropriately sized vertex set. The distinction between abstract and geometric complexes isn't that strict, and we will often identify an abstract simplicial complex with a specific geometric realization. The main advantage of abstract complexes is that they are simple to define.

Definition 2.9. For two simplicial complexes K and L, a **simplicial mapping** from K to L is a function $f : V(K) \to V(L)$ such that $f(\sigma) \in L$ for all $\sigma \in K$. In other words, it's a function on the vertices which preserves the simplex status of a collection of vertices.

There is a way to go from a map on the vertices of a simplicial complex to a map on the interior of every simplex:

Definition 2.10. For two geometric simplicial complexes Δ_1 and Δ_2 and a simplicial map $f : V(\Delta_1) \to V(\Delta_2)$, the **affine extension** of f is the extension of f onto $\|\Delta_1\|$,

$$\|f\| : \|\Delta_1\| \to \|\Delta_2\|$$

given by

$$\|f\| (\boldsymbol{x}) \coloneqq \sum_{i=0}^{n} \alpha_i f(\boldsymbol{v}_i)$$

where $\boldsymbol{x}$ is in the simplex $[\boldsymbol{v}_0, \dots, \boldsymbol{v}_n]$ with coordinates $\alpha_0, \dots, \alpha_n$

This is well-defined, as a point on two simplices must be on a common face, so any coordinate it has for a vertex which isn't shared must be zero.

The affine extension of a simplicial mapping is continuous, because it's continuous on each of the relative interiors of simplices in the complex, and around every vertex you can find a neighborhood where the coordinates other than the one for the vertex in question are arbitrarily small.

2.2. **Homotopy and Homology.** We're now going to move into the algebraic topology necessary for our proofs of the Borsuk-Ulam Theorem and Brouwer's Fixed Point Theorem. This is meant mostly to collect the necessary information to be able to read the proofs; for a more complete understanding, please consult Hatcher's textbook [2].

We first define homotopy. This gives a way of deforming one shape into another one which only preserves some very basic properties.

Definition 2.11. Two continuous functions $f_0, f_1 : X \to Y$ are **homotopic**, written $f_0 \simeq f_1$, if there is a continuous function $F : X \times I \to Y$ with $F(x, 0) = f_0(x)$ and $F(x, 1) = f_1(x)$ for all x. The function F is called a **homotopy**.

Definition 2.12. Two spaces X and Y are **homotopy equivalent** if there are continuous functions $f : X \to Y$ and $g : Y \to X$ such that $f \circ g \simeq \mathrm{id}_Y$ and $g \circ f \simeq \mathrm{id}_X$

One type of homotopy equivalence is a deformation retraction:

Definition 2.13. A subspace A of X is a **deformation retract** of the space X if there is a continuous function $r : X \to A$ which is the identity on A, and a homotopy $F : X \times I \to X$ between r and id_X which is always the identity on A.

We now move on to homology groups. The homology groups of a space are a way to count holes in a space, so they are meant to represent all the loops in the space, considering two loops the same if one can be turned into the other without crossing a hole. We'll specifically deal with "singular homology".

Definition 2.14. An **oriented simplex** is a simplex where the order of the vertices is taken into account. Every oriented simplex which is derived from $[\boldsymbol{v}_0, \dots, \boldsymbol{v}_n]$ by performing an even permutation on the vertices is considered the same oriented simplex. There are thus two orientations for every set of vertices. We'll write σ and $-\sigma$ for the two orientations.

Definition 2.15. A **singular n-simplex** in a space X is a continuous function $\sigma : \Delta^n \to X$. We'll think of σ as a simplex which was placed in X.

If you reverse the orientation of Δ^n, call the resulting singular simplex $-\sigma$.

Definition 2.16. An n-chain is defined to be a finite formal sum of singular n-simplices. We add an n-chain "0", and define $\sigma + (-\sigma) = 0$.

This definition has every n-chain written as a sum of simplices, with integer coefficients. We can actually define n-chains to have coefficients in any abelian group, not just the integers.

To be formal, we let $C_n(X; G)$ be the abelian group with elements of the form $\sum_i g_i\sigma_i$, where $g_i \in G$ and $\sigma_i : \Delta^n \to X$. We also give the group left and right distributive laws. We call this **the group of n-chains in X with coefficients in G.**

We define a boundary map for the group of n-chains. The reason for this is that it will shortly be used to define a loop (called a "cycle").

Definition 2.17. The **boundary map** $\partial_n : C_n(X; G) \to C_{n-1}(X; G)$ is defined by first setting its value on each singular n-simplex to be

$$\partial_n(\sigma) \coloneqq \sum_{i=0}^{n} (-1)^i \sigma|[\boldsymbol{v}_0, \dots, \hat{\boldsymbol{v}}_i, \dots, \boldsymbol{v}_n]$$

where the hat over the vertex means we remove it from the simplex. Then we extend the boundary map linearly.

We will generally write the application of or composition with ∂ without parentheses or a "$\circ$" symbol. So we'll write these like $\partial\sigma$ to mean $\partial(\sigma)$, ∂f to mean $\partial \circ f$, and $f\partial$ to mean $f \circ \partial$.

There are two important features of the boundary map which you can verify:

(1) For all n, $\partial_n\partial_{n+1} = 0$. (Sometimes written $\partial^2 = 0$, as the subscripts are usually omitted.)
(2) In the cases we can easily visualize ($n = 1, 2$), $\partial\xi = 0$ if and only if ξ forms a loop (taking orientation into account, as well).

The second property isn't rigorous, as we haven't defined loops. Instead, it gives us inspiration for how we can define them:

Definition 2.18. A **cycle** is an n-chain ξ such that $\partial\xi = 0$.

The equation $\partial^2 = 0$ can be rephrased as "every boundary is a cycle".

We are ready to define the homology groups. One way to sensibly get from one cycle to another that would make them "the same" is to add a boundary. This is how we define the homology groups:

Definition 2.19. For a space X and abelian group G, the **n-th homology group with coefficients in G** is

$$H_n(X;G) \coloneqq \operatorname{Ker}\partial_n / \operatorname{Im}\partial_{n+1}$$

We let $[\xi]$ be the set $\{\xi + \partial\alpha \ : \ \alpha \in C_{n+1}(X;G)\}$, the **homology class** of ξ.

The most common coefficients are $\mathbb{Z}$. Unless otherwise stated, if we don't write the coefficients, they are $\mathbb{Z}$.

Homology as a concept isn't tied to chains of singular simplices. All you need is a sequence of abelian groups $\{C_n\}$, and a family of maps $\partial_n : C_n \to C_{n-1}$ with $\partial^2 = 0$. Such a structure is called a **chain complex**, and the homology groups of the complex can be defined as $\operatorname{Ker}\partial / \operatorname{Im}\partial$.

Definition 2.20. Until now, we've assumed there are no negative-dimensional chains, so the chain complex we used to define the homology groups ends with C_0, and then the trivial group.

However, if we define $C_{-1}(X;G)$ as the group of chains deriving from the singular empty simplex (of which there is exactly one), we get the **reduced homology groups**, $\tilde{H}_n(X;G)$.

The boundary map $C_0(X;G) \to C_{-1}(X;G)$, denoted ε, is

$$\varepsilon\left(\sum_i g_i\sigma_i\right) \coloneqq \sum_i g_i$$

We can also just think of ε as a homomorphism $C_0(X;G) \to G$. This homomorphism contains $\operatorname{Im}\partial_1$ in its kernel, so it induces a homomorphism $H_0(X;G) \to G$, with kernel $\tilde{H}_0(X;G)$. Therefore, $H_0(X;G) \approx \tilde{H}_0(X;G) \oplus G$. This is the only dimension where the reduced homology group is different.

We can now introduce the results we'll need in this paper.

Theorem 2.21. *If X is a point, then $\tilde{H}_n(X;G) \approx 0$ for all n.*

Proof. For each n, there is only one singular n-simplex, σ_n. So, $C_n(X;G) \approx G$. Calculating the boundary of σ_n gives $\partial\sigma_n = \sum_{i=0}^n (-1)^i \sigma_{n-1}$, so the boundary map is zero for odd n, and an isomorphism for even n. The chain complex looks like

$$\cdots \xrightarrow{0} G \xrightarrow{\approx} G \xrightarrow{0} G \xrightarrow{\approx} G \longrightarrow 0$$

At every part of the chain complex, $\operatorname{Ker}\partial / \operatorname{Im}\partial \approx 0$. ∎

Theorem 2.22. *A continuous map $f : X \to Y$ induces homomorphisms $f_\sharp : C_n(X;G) \to C_n(Y;G)$ and $f_* : H_n(X;G) \to H_n(Y;G)$.*

Proof. We let $f_\sharp(\sigma) \coloneqq f \circ \sigma$, and extend linearly. Because we extended linearly, we get that the following diagram **commutes**:

$$\begin{tikzcd}
C_{n+1}(X;G) \arrow[r, "\partial"] \arrow[d, "f_\sharp"] & C_n(X;G) \arrow[r, "\partial"] \arrow[d, "f_\sharp"] & C_{n-1}(X;G) \arrow[d, "f_\sharp"] \\
C_{n+1}(Y;G) \arrow[r, "\partial"] & C_n(Y;G) \arrow[r, "\partial"] & C_{n-1}(Y;G)
\end{tikzcd}$$

A **commutative diagram** is one where any two sequences of arrows which start and end in the same places correspond to two compositions of functions which are equal. The commutativity of the above diagram means $f_\sharp\partial = \partial f_\sharp$.

Because of this identity, $f_\sharp$ sends cycles to cycles and boundaries to boundaries, and hence induces a homomorphism f_* on the homology groups by $f_*[\alpha] = [f_\sharp(\alpha)]$. □

Definition 2.23. A map from one chain complex to another which commutes with the boundary map is called a **chain map**.

So what we really just showed is that chain maps induce homomorphisms on the homology groups, and continuous function induce chain maps.

Theorem 2.24. *If $f, g : X \to Y$ are homotopic, then they induce the same homomorphism on the homology groups.*

We won't give all the details here. For the complete proof (with pictures), see [2], page 112.

Proof. First, you form a certain subdivision of $\Delta^n \times I$ into simplices. Specifically, if $\Delta^n \times \{0\} = [v_0, v_1, \ldots, v_n]$ and $\Delta^n \times \{1\} = [w_0, w_1, \ldots, w_n]$, then we subdivide into simplices of the form

$$\sigma_{n-i} = [v_0, \ldots, v_i, w_i, w_{i+1}, \ldots, w_n]$$

If $F(x,t)$ is a homotopy from f to g, and σ is a singular n-simplex in X, we can define the composition $F \circ (\sigma \times \mathrm{id}) : \Delta^n \times I \to X \times I \to Y$.

All this is to define the "prism operator" $P : C_n(X) \to C_{n+1}(Y)$. Let

$$P(\sigma) \coloneqq \sum_{i=0}^{n} (-1)^i F \circ (\sigma \times \mathrm{id})|\sigma_i$$

This function has an important property:

$$\partial P = g_\sharp - f_\sharp - P\partial$$

Given this property, if α is a cycle, then

$$\begin{aligned} g_\sharp(\alpha) - f_\sharp(\alpha) &= \partial P(\alpha) - P\partial(\alpha) \\ &= \partial P(\alpha) \end{aligned}$$

So the induced homomorphisms g_* and f_* send the homology class of α to the same homology class. □

Since $(f \circ g)_* = f_* \circ g_*$, this theorem also implies that homotopy equivalent spaces has isomorphic homology groups.

We introduce an important algebraic tool.

Definition 2.25. An **exact sequence** is a sequence of groups $\{A_n\}$ and homomorphisms $\alpha_n : A_n \to A_{n-1}$ such that $\operatorname{Ker} \alpha_n = \operatorname{Im} \alpha_{n+1}$.

An exact sequence of the form

$$0 \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \longrightarrow 0$$

is called a **short exact sequence**. Note that α must be injective, and β must be surjective.

This is an important result, which we won't be able to prove entirely. We'll prove one of the intermediate results, and present the general sketch.

Theorem 2.26. *If A is a deformation retract of an open neighborhood in X, then there is an exact sequence*

$$\cdots \longrightarrow \tilde{H}_n(A) \xrightarrow{i_*} \tilde{H}_n(X) \xrightarrow{j_*} \tilde{H}_n(X/A) \xrightarrow{\partial} \tilde{H}_{n-1}(A) \longrightarrow \cdots \longrightarrow \tilde{H}_0(X/A) \longrightarrow 0$$

where i is the inclusion map, j is the quotient map, and the coefficients are in any group G, for some homomorphism ∂.

We will show this intermediate result, which gets you to a long exact sequence of homology groups, as we'll use this result separately:

Theorem 2.27. *If A_n, B_n, and C_n are chain complexes, with chain maps i, j such that the following sequence is exact for all n:*

$$0 \longrightarrow A_n \xrightarrow{i} B_n \xrightarrow{j} C_n \longrightarrow 0$$

then there is a map ∂ giving a long exact sequence of homology groups:

$$\cdots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(B) \xrightarrow{j_*} H_n(C) \xrightarrow{\partial} H_{n-1}(A) \longrightarrow \cdots$$

Proof. Note that the exactness of the given sequence implies that i is one-to-one, j is onto, and we have $j \circ i = 0$.

We will define $\partial : H_n(C) \to H_{n-1}(A)$. Let $c \in C_n$ be a cycle. Because j is onto, there is $b \in B_n$ such that $j(b) = c$. Taking the boundary of b, we get something in $\operatorname{Ker} j$, because $j(\partial b) = \partial j(b) = \partial c = 0$. By exactness, we have that ∂b is in $\operatorname{Im} i$, so there is some $a \in A_{n-1}$ with $i(a) = \partial b$. This chain a is a cycle, because $i(\partial a) = \partial i(a) = \partial\partial b = 0$, and the homomorphism i is one-to-one. So we can let $\partial[c] \coloneqq [a]$.

To see that ∂ is well-defined, we need to show that at any point where choice is involved, we always send a given homology class to the same homology class.

- Once c and b are chosen, ∂b and thus a are unique, by the injectivity of i.
- If there were a different choice for b, that is, a b' such that $c = j(b')$, then we'd have $b' - b \in \operatorname{Ker} j = \operatorname{Im} i$. Therefore, there would be some $a' \in A_n$ with $i(a') = b' - b$, so that $b' = b + i(a')$. We'd then have $\partial b' = \partial(b + i(a')) = \partial b + i(\partial a') = i(a) + i(\partial a') = i(a + \partial a')$, so we'd send $[c]$ to $[a + \partial a']$. But this is the same homology class as $[a]$.
- A different choice from the homology class of c would be of the form $c + \partial c'$. By surjectivity, there is some $b' \in B_{n+1}$ such that $c' = j(b')$, so $c + \partial c' = j(b + \partial b')$, and b is replaced with $b + \partial b'$. But this has the same boundary as b, so the rest is unchanged.

We also need to check that ∂ is a homomorphism. Suppose $\partial[c_1] = [a_1]$ and $\partial[c_2] = [a_2]$, via b_1 and b_2. Then $c_1 + c_2 = j(b_1 + b_2)$, and $\partial(b_1 + b_2) = \partial b_1 + \partial b_2 = i(a_1) + i(a_2) = i(a_1 + a_2)$. So $\partial([c_1] + [c_2]) = [a_1] + [a_2]$, as desired.

Lastly, we need to verify that the produced sequence is exact. We'll work through the inclusions:

- $\operatorname{Im} i_* \subseteq \operatorname{Ker} j_*$: We're given $j \circ i = 0$, so $j_* \circ i_* = 0$.
- $\operatorname{Im} j_* \subseteq \operatorname{Ker} \partial$: If $[c] \in \operatorname{Im} j_*$, then b in the construction of ∂ is a cycle, so $\partial b = 0$, and hence $\partial[c] = 0$
- $\operatorname{Im} \partial \subseteq \operatorname{Ker} i_*$: In the notation of the construction, $i_*\partial[c] = i_*[a] = [i(a)] = [\partial b] = 0$.

- $\operatorname{Ker} j_* \subseteq \operatorname{Im} i_*$: If $[b] \in \operatorname{Ker} j_*$, then $j(b) = \partial c'$ for some $c' \in C_{n+1}$. Because j is onto, there is some $b' \in B_{n+1}$ such that $j(b') = c'$. The chain $b - \partial b'$ is in the same homology class as b. It is also in $\operatorname{Ker} j$, as $j(b - \partial b') = j(b) - \partial j(b') = \partial c' - \partial c' = 0$. Because $\operatorname{Ker} j = \operatorname{Im} i$, there is some $a \in A_n$ such that $i(a) = b - \partial b'$. This a is a cycle, as $i(\partial a) = \partial i(a) = \partial(b - \partial b') = \partial b = 0$. Therefore, $i_*[a] = [b - \partial b'] = [b]$.
- $\operatorname{Ker} \partial \subseteq \operatorname{Im} j_*$: In the notation of the construction, if $[c] \in \operatorname{Ker} \partial$, then $a = \partial a'$ for some $a' \in A_n$. Consider $b - i(a')$. This is a cycle, because $\partial(b - i(a')) = \partial b - i(\partial a') = i(a) - i(a) = 0$. The function j sends this cycle to c, because $j(b - i(a')) = j(b) - j(i(a')) = c$. So, $j_*[b - i(a')] = [c]$.
- $\operatorname{Ker} i_* \subseteq \operatorname{Im} \partial$: If $[a] \in \operatorname{Ker} i_*$, then $i(a) = \partial b$, for some $b \in B_n$. We have $\partial j(b) = j(\partial b) = j(i(a)) = 0$, so $j(b)$ is a cycle. By the definition of ∂, we have $\partial[j(b)] = [a]$.

□

If we have the elements necessary to apply this theorem, we say we have a **short exact sequence of chain complexes**.

It's worth checking that we actually can't apply this to immediately get the sequence in Theorem 2.26, as the quotient map from X to X/A doesn't induce a chain map which gives an exact sequence. What we can instead do to prove Theorem 2.26 is define a new chain complex $C_n(X, A) \coloneqq C_n(X)/C_n(A)$ and homology groups $H_n(X, A)$, with which we can apply this result. Then we use a different result, called the "Excision Theorem," to show that $H_n(X, A) \approx \tilde{H}_n(X/A)$ when A is a deformation retract of an open neighborhood in X. (For the full proof, see [2], pages 114-124.)

We can say a little bit more about the sequence in Theorem 2.27. Namely,

Theorem 2.28. *Suppose we have a short exact sequence of chain complexes as in Theorem 2.27, and another short exact sequence of chain complexes (denoted by A'_n, B'_n, and C'_n). If there are families of functions α, β, and γ such that the diagram*

$$\begin{tikzcd}
0 \arrow[r] & A_n \arrow[r, "i"] \arrow[d, "\alpha"] & B_n \arrow[r, "j"] \arrow[d, "\beta"] & C_n \arrow[r] \arrow[d, "\gamma"] & 0 \\
0 \arrow[r] & A'_n \arrow[r, "i'"] & B'_n \arrow[r, "j'"] & C'_n \arrow[r] & 0
\end{tikzcd}$$

commutes for all n, then the induced homology diagram also commutes:

$$\begin{tikzcd}
\cdots \arrow[r] & H_n(A) \arrow[r, "i_*"] \arrow[d, "\alpha_*"] & H_n(B) \arrow[r, "j_*"] \arrow[d, "\beta_*"] & H_n(C) \arrow[r, "\partial"] \arrow[d, "\gamma_*"] & H_{n-1}(A) \arrow[r] \arrow[d, "\alpha_*"] & \cdots \\
\cdots \arrow[r] & H_n(A') \arrow[r, "i'_*"] & H_n(B') \arrow[r, "j'_*"] & H_n(C') \arrow[r, "\partial"] & H_{n-1}(A') \arrow[r] & \cdots
\end{tikzcd}$$

Proof. We can go through the homology diagram shown, showing each square commutes.

For the left square, we're given that $\beta \circ i = i' \circ \alpha$, so we have $\beta_* \circ i_* = i'_* \circ \alpha_*$

Similarly, for the middle square, we're given $\gamma \circ j = j' \circ \beta$, so we have $\gamma_* \circ j_* = j'_* \circ \beta_*$.

For the right square, recall the definition of ∂. Suppose $[a] \coloneqq \partial[c]$ and $j(b) \coloneqq c$, with $i(a) = \partial b$. Then $\gamma(c) = \gamma(j(b)) = j'(\beta(b))$, and $i'(\alpha(a)) = \beta(i(a)) = \beta(\partial(b)) = \partial(\beta(b))$. So, according to the definition of ∂ in the A'_n-B'_n-C'_n short exact sequence, we have $\partial[\gamma(c)] = [\alpha(a)]$. In other words, $\partial\gamma_*[c] = \alpha_*\partial[c]$. $\square$

3. Borsuk-Ulam and Brouwer's Fixed Point Theorem

First, we'll explain the statements of the two classic results, and how the Borsuk-Ulam Theorem implies Brouwer's Fixed Point Theorem. Then we'll present a topological proof for both theorems.

3.1. The Borsuk-Ulam Theorem.

The usual version of the Borsuk-Ulam Theorem says that any continuous map from S^n to $\mathbb{R}^n$ sends a pair of **antipodal points** (points on the sphere that are diametrically opposite) to the same point:

Theorem 3.1. *Let $f : S^n \to \mathbb{R}^n$ be a continuous function. Then there is a point $\boldsymbol{x} \in S^n$ such that $f(\boldsymbol{x}) = f(-\boldsymbol{x})$.*

The Borsuk-Ulam Theorem has several similar, equivalent formulations. We only state one here, but a longer list can be found in [1]: (B^n denotes the closed n-dimensional unit ball)

Theorem 3.2. *There is no continuous function $f : B^n \to S^{n-1}$ which is* ***odd on the boundary*** *(that is, $f(-\boldsymbol{x}) = -f(\boldsymbol{x})$ for all points $\boldsymbol{x}$ on $\partial B^n = S^{n-1}$).*

Proof of Theorem 3.2 from Theorem 3.1. Let $\pi : S^n \to B^n$ be the projection

$$(x_1, \ldots, x_n, x_{n+1}) \longmapsto (x_1, \ldots, x_n)$$

Suppose there is a continuous function $f : B^n \to S^{n-1}$ which is antipodal on the boundary. Then $g = f \circ \pi$ is a continuous function from S^n to S^{n-1} which is odd. On S^{n-1}, which doesn't contain the origin, this implies that $g(\boldsymbol{x}) \neq g(-\boldsymbol{x})$ for all $\boldsymbol{x} \in S^n$. As S^{n-1} is a subset of $\mathbb{R}^n$, this contradicts Theorem 3.1, and hence f can't exist. $\square$

Proof of Theorem 3.1 from Theorem 3.2. Suppose there is a continuous function $f : S^n \to \mathbb{R}^n$ with $f(\boldsymbol{x}) \neq f(-\boldsymbol{x})$ for all points $\boldsymbol{x} \in S^n$.

We can then define a function $g : S^n \to S^{n-1}$ by $g(\boldsymbol{x}) \coloneqq \frac{f(\boldsymbol{x}) - f(-\boldsymbol{x})}{\|f(\boldsymbol{x}) - f(-\boldsymbol{x})\|}$. Note that g is odd.

Let $\pi^{-1} : B^n \to S^n$ be the inverse of the projection of the upper hemisphere:

$$\boldsymbol{x} \longmapsto \left(\boldsymbol{x}, \sqrt{1 - \|\boldsymbol{x}\|^2}\right)$$

Then $h \coloneqq g \circ \pi^{-1}$ is a continuous function from B^n to S^{n-1} which is odd on the boundary, contradicting Theorem 3.2. $\square$

3.2. Brouwer's Fixed Point Theorem.

Brouwer's Fixed Point Theorem in general guarantees a fixed point for any continuous map from a compact, convex subset of $\mathbb{R}^n$ to itself. It is sufficient to state it for the n-dimensional closed ball:

Theorem 3.3. *Every continuous function $f : B^n \to B^n$, has a* ***fixed point****. That is, a point $\boldsymbol{x} \in B^n$ with $f(\boldsymbol{x}) = \boldsymbol{x}$.*

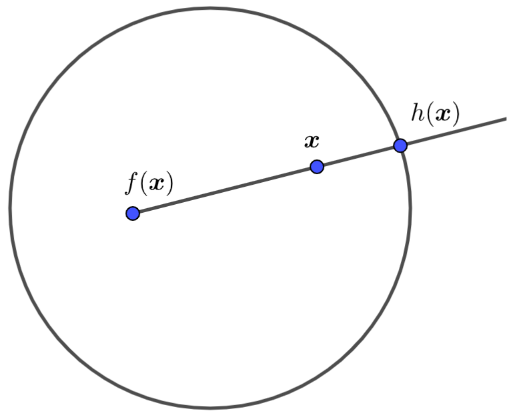

FIGURE 1. Demonstration of retraction when $n = 2$

3.3. **Proof that the Borsuk-Ulam Theorem Implies Brouwer's Fixed Point Theorem.** The proof relies upon a construction called a "retraction". If we suppose that Brouwer's Fixed Point Theorem is false for some f, then we can draw a ray from $f(\boldsymbol{x})$ through $\boldsymbol{x}$ for all $\boldsymbol{x} \in B^n$, and let $h(\boldsymbol{x})$ be where that ray intersects the boundary, S^{n-1}. By inspection, h is continuous. (Consider how h acts on a neighborhood around a point, and use the continuity of f.) Moreover, h is the identity on S^{n-1}. These two properties are what makes h a **retraction** – a continuous function from a space X onto a subspace A which is the identity on A. This construction will also be used in the proof of Brouwer's Fixed Point Theorem given later, which does not use the Borsuk-Ulam Theorem.

So, if Brouwer's Fixed Point Theorem is false, then we can construct a retraction of B^n onto S^{n-1}. However, such a function would be odd on the boundary (because it's the identity on the boundary). If the Borsuk-Ulam Theorem is true, no such function exists, so Brouwer's Fixed Point Theorem must be true.

3.4. **Proof of Brouwer's Fixed Point Theorem.** We've just laid the groundwork for the general proof of Brouwer's Fixed Point Theorem. The strategy will be to show that there cannot be a retraction of B^n onto S^{n-1}.

This calculation will be necessary for this proof, and the proof of the Borsuk-Ulam Theorem.

Lemma 3.4. $\tilde{H}_k(S^n; G) = \begin{cases} G, & k = n \\ 0, & \textit{otherwise} \end{cases}$

Proof. It is a fact from topology that S^n is homeomorphic to B^n/S^{n-1}. So by Theorem 2.26, there is a long exact sequence:

$$\cdots \longrightarrow \tilde{H}_k(S^{n-1}) \longrightarrow \tilde{H}_k(B^n) \longrightarrow \tilde{H}_k(S^n) \longrightarrow \tilde{H}_{k-1}(S^{n-1}) \longrightarrow \cdots$$

(The coefficients are in G.) The space B^n deformation retracts to a point, by the homotopy $F(\boldsymbol{x}, t) = t\boldsymbol{x}$. So by Theorems 2.21 and 2.24, $\tilde{H}_k(B^n; G) \approx 0$ for all k

and n. Therefore, there is an exact sequence

$$0 \longrightarrow \tilde{H}_k(S^n) \longrightarrow \tilde{H}_{k-1}(S^{n-1}) \longrightarrow 0$$

for all positive k and n, and for all positive n there is an exact sequence

$$0 \longrightarrow \tilde{H}_0(S^n) \longrightarrow 0$$

So $\tilde{H}_k(S^n; G) \approx \tilde{H}_{k-1}(S^{n-1}; G)$, and $\tilde{H}_0(S^n; G) \approx 0$, for all positive k and n.

S^0 is two points, so in general there are two singular k-simplices, α_k and β_k. Like in the proof of Theorem 2.21, $\partial(\alpha_k)$ is 0 for odd k and α_{k-1} for even k, and similarly for β_k. Moreover, $C_k(S^0; G) \approx G \oplus G$, so the chain complex is

$$\cdots \xrightarrow{0} G \oplus G \xrightarrow{\approx} G \oplus G \xrightarrow{0} G \oplus G \xrightarrow{\varepsilon} G \longrightarrow 0$$

where $\operatorname{Ker} \varepsilon = \{(g, -g) \ : \ g \in G\}$. Hence $\tilde{H}_k(S^0; G)$ is isomorphic to G for $k = 0$ and trivial otherwise.

By induction, we get our result. ☐

We can use this calculation to show that S^{n-1} is not a retract of B^n.

Proof. Suppose $r : B^n \to S^{n-1}$ is a retraction. Then if $i : S^{n-1} \to B^n$ is the inclusion map, we have $r \circ i = \mathrm{id}$. Passing to homology groups, then we have $r_* \circ i_* = \mathrm{id}_*$, which is the identity on $\tilde{H}_{n-1}(S^{n-1})$. But $\tilde{H}_{n-1}(B^n)$ is trivial, so i_* and r_* are both trivial maps. However, $\tilde{H}_{n-1}(S^{n-1})$ is not trivial, so we have a contradiction.

Brouwer's Fixed Point Theorem follows from the construction of a retraction we did earlier. ☐

3.5. **Proof of the Borsuk-Ulam Theorem.** The proof of the Borsuk-Ulam Theorem relies upon a lemma about the "degree" of a function. The degree of a function is, intuitively, how many times it "wraps around" the origin. To uncerstand the connection to polynomial degree, look at a complex polynomial, like $f(z) = z^k$. Any point $e^{i\theta}$ on the unit circle has exactly k preimages under f. As this sort of property generalizes more easily than the degree of a polynomial, it's used to define degree for all continuous functions. Homology gives a very clean way to phrase it:

Definition 3.5. For a continuous function $f : S^n \to S^n$, the induced homomorphism $f_* : H_n(S^n) \to H_n(S^n)$ further induces a homomorphism f'_* from $\mathbb{Z}$ to $\mathbb{Z}$. Hence, there is an integer d such that $f'_*(\alpha) = d\alpha$. The **degree** of f is this integer d.

We need to perform the following calculation for our lemma. The reason will become clear shortly.

Lemma 3.6. *$H_k(\mathbb{R}\mathrm{P}^n; G) \approx 0$ for all $k > n$.*

Proof. It's a fact from topology that S^n is homeomorphic to $\mathbb{R}\mathrm{P}^n/\mathbb{R}\mathrm{P}^{n-1}$, so by Theorem 2.26 we get an exact sequence

$$\cdots \longrightarrow H_k(\mathbb{R}\mathrm{P}^{n-1}) \longrightarrow H_k(\mathbb{R}\mathrm{P}^n) \longrightarrow H_k(S^n) \longrightarrow H_{k-1}(\mathbb{R}\mathrm{P}^{n-1}) \longrightarrow \cdots$$

with coefficients in G. For $k > n$, we get the following exact sequence by Lemma 3.4:

$$0 \longrightarrow H_k(\mathbb{R}\mathrm{P}^{n-1}) \longrightarrow H_k(\mathbb{R}\mathrm{P}^n) \longrightarrow 0$$

When $n = 0$, we have $\mathbb{R}\mathrm{P}^n$ is a point, so our lemma is true by Theorem 2.21. By induction with this short exact sequence, we get our result for all n. ☐

The lemma is this:

Lemma 3.7. *An odd map $f : S^n \to S^n$ has odd degree.*

Proof. The goal of the proof is to show that $f_* : H_n(S^n; \mathbb{Z}_2) \to H_n(S^n; \mathbb{Z}_2)$ is an isomorphism, which will imply that f has odd degree. To that end, it's useful to produce a long exact sequence with the homology groups of S^n and $\mathbb{R}\mathrm{P}^n$ with coefficients in $\mathbb{Z}_2$.

Let $p : S^n \to \mathbb{R}\mathrm{P}^n$ be the quotient map $p(x) = [x] = \{x, -x\}$. We get an associated chain map $p_\sharp : C_n(S^n; \mathbb{Z}_2) \to C_n(\mathbb{R}\mathrm{P}^n; \mathbb{Z}_2)$. This map is onto.

Every singular n-simplex $\sigma : \Delta^n \to \mathbb{R}\mathrm{P}^n$ has exactly two singular n-simplices $\tilde{\sigma}_1, \tilde{\sigma}_2 : \Delta^n \to S^n$ which get sent to it by $p_\sharp$. Because our coefficients are in $\mathbb{Z}_2$, the kernel of $p_\sharp$ is generated by all such $\tilde{\sigma}_1 + \tilde{\sigma}_2$.

If we define $\tau(\sigma) \coloneqq \tilde{\sigma}_1 + \tilde{\sigma}_2$ and extend to chains linearly, then $\operatorname{Ker} p_\sharp = \operatorname{Im} \tau$, so we get the following short exact sequence:

$$0 \longrightarrow C_n(\mathbb{R}\mathrm{P}^n; \mathbb{Z}_2) \xrightarrow{\tau} C_n(S^n; \mathbb{Z}_2) \xrightarrow{p_\sharp} C_n(\mathbb{R}\mathrm{P}^n; \mathbb{Z}_2) \longrightarrow 0$$

Both τ and $p_\sharp$ commute with the boundary map, so this extends to a short exact sequence of chain complexes, and by Theorem 2.27, this gives us a long exact sequence of homology groups. For convenience, we'll abbreviate $\mathbb{R}\mathrm{P}^n$ as P^n and leave the $\mathbb{Z}_2$ coefficients implicit:

$$\cdots \longrightarrow H_k(P^n) \xrightarrow{\tau_*} H_k(S^n) \xrightarrow{p_*} H_k(P^n) \longrightarrow H_{k-1}(P^n) \longrightarrow \cdots$$

By Lemmas 3.4 and 3.6, we know many of the groups in this sequence:

$$0 \longrightarrow H_n(P^n) \longrightarrow \mathbb{Z}_2 \longrightarrow H_n(P^n) \longrightarrow H_{n-1}(P^n) \longrightarrow 0 \longrightarrow \cdots$$

$$\cdots \longrightarrow 0 \longrightarrow H_i(P^n) \longrightarrow H_{i-1}(P^n) \longrightarrow 0 \longrightarrow \cdots$$

$$\cdots \longrightarrow 0 \longrightarrow H_1(P^n) \longrightarrow H_0(P^n) \longrightarrow \mathbb{Z}_2 \longrightarrow H_0(P^n) \longrightarrow 0$$

The middle row tells us that $H_i(P^n) \approx H_{i-1}(P^n)$, for $1 < i < n$.

From the bottom row, we know that $H_0(P^n)$ is not trivial, as that would imply $\mathbb{Z}_2$ is trivial. We also know from the bottom row that $|H_0(P^n)| \leq |\mathbb{Z}_2|$ because there's a surjection from $\mathbb{Z}_2$ to $H_0(P^n)$. So $|H_0(P^n)| = 2$, and hence $H_0(P^n) \approx \mathbb{Z}_2$.

From this, we infer that the map from $H_0(P^n)$ to $\mathbb{Z}_2$ is trivial, so $H_1(P^n) \approx H_0(P^n)$. By induction, $H_i(P^n) \approx \mathbb{Z}_2$ for $i < n$.

From the top row, $|H_{n-1}(P^n)| \leq |H_n(P^n)| \leq |\mathbb{Z}_2|$. But then by the calculations we've done already, $|H_n(P^n)| = 2$, and hence $H_n(P^n) \approx \mathbb{Z}_2$.

(Note that the case of $n = 1$ is slightly different. We still get that the map from $H_0(P^1)$ to $\mathbb{Z}_2$ is trivial, but this only lets us deduce that $|H_0(P^1)| \leq |H_1(P^1)|$. However, we also still get that $|H_1(P^1)| \leq |\mathbb{Z}_2|$, so the result is the same. The case of $n = 0$ is also different, but it's a very simple exact sequence in that case.)

Because we know what all the groups are, up to isomorphism, we can deduce that the maps must behave like this:

$$0 \longrightarrow H_n(P^n) \xrightarrow{\approx} H_n(S^n) \xrightarrow{0} H_n(P^n) \xrightarrow{\approx} H_{n-1}(P^n) \longrightarrow 0 \longrightarrow \cdots$$

$$\cdots \longrightarrow 0 \longrightarrow H_i(P^n) \xrightarrow{\approx} H_{i-1}(P^n) \longrightarrow 0 \longrightarrow \cdots$$

$$\cdots \longrightarrow 0 \longrightarrow H_1(P^n) \xrightarrow{\approx} H_0(P^n) \xrightarrow{0} H_0(S^n) \xrightarrow{\approx} H_0(P^n) \longrightarrow 0$$

All of this was to set up the above diagram. Now, if $f : S^n \to S^n$ is odd, then it induces a map $\bar{f} : \mathbb{R}\mathrm{P}^n \to \mathbb{R}\mathrm{P}^n$ defined by $\bar{f} \circ p = p \circ f$. (The oddness of f is necessary here.) Let $f_\sharp$ and $\bar{f}_\sharp$ be the induced chain maps.

This diagram commutes for all i:

$$\begin{tikzcd}
0 \arrow[r] & C_i(P^n) \arrow[r, "\tau"] \arrow[d, "\bar{f}_\sharp"] & C_i(S^n) \arrow[r, "p_\sharp"] \arrow[d, "f_\sharp"] & C_i(P^n) \arrow[r] \arrow[d, "\bar{f}_\sharp"] & 0 \\
0 \arrow[r] & C_i(P^n) \arrow[r, "\tau"] & C_i(S^n) \arrow[r, "p_\sharp"] & C_i(P^n) \arrow[r] & 0
\end{tikzcd}$$

The right square commutes by the definition of $\bar{f}$. To show the left square commutes, let a simplex $\sigma \in C_i(P^n)$ have two simplices $\tilde{\sigma}_1, \tilde{\sigma}_2 \in C_i(S^n)$ which get sent to it by $p_\sharp$. Note that, because f is odd, $f_\sharp(\tilde{\sigma}_1) \neq f_\sharp(\tilde{\sigma}_2)$. We have $p_\sharp(f_\sharp\tilde{\sigma}_1) = \bar{f}_\sharp(p_\sharp\tilde{\sigma}_1) = \bar{f}_\sharp\sigma$, and similarly for $\tilde{\sigma}_2$, so $\tau(\bar{f}_\sharp(\sigma)) = f_\sharp(\tilde{\sigma}_1) + f_\sharp(\tilde{\sigma}_2) = f_\sharp(\tau(\sigma))$.

By Theorem 2.28, this extends to a commutative diagram of the homology sequence we've just computed. We won't show the whole diagram, but it contains these commutative squares for all i:

$$\begin{tikzcd}
H_i(P^n) \arrow[r, "\approx"] \arrow[d, "\bar{f}_*"] & H_{i-1}(P^n) \arrow[d, "\bar{f}_*"] \\
H_i(P^n) \arrow[r, "\approx"] & H_{i-1}(P^n)
\end{tikzcd}
\qquad
\begin{tikzcd}
H_n(P^n) \arrow[r, "\approx"] \arrow[d, "\bar{f}_*"] & H_n(S^n) \arrow[d, "f_*"] \\
H_n(P^n) \arrow[r, "\approx"] & H_n(S^n)
\end{tikzcd}$$

We'll show $\bar{f}_* : H_0(P^n) \to H_0(P^n)$ is an isomorphism, and use induction to show that $f_* : H_n(S^n) \to H_n(S^n)$ is an isomorphism.

Suppose σ_1, σ_2 are singular 0-simplices in $\mathbb{R}\mathrm{P}^n$. Because $\mathbb{R}\mathrm{P}^n$ is path-connected, there is a singular 1-simplex σ so that $\partial\sigma = \sigma_1 - \sigma_2$. So in $C_0(P^n)$, which is in $\mathbb{Z}_2$ coefficients, $\sigma_1 + \sigma_2$ is a boundary. Therefore, any chain in $C_0(P^n)$ can have its terms grouped in pairs to get $c = a + \partial b$, where a is zero or a singular 0-simplex. We already know that $H_0(P^n) \approx \mathbb{Z}_2$, and any two chains of the form $a + \partial b$, $a \neq 0$, differ by a boundary, so the homology classes of $H_0(\mathbb{R}\mathrm{P}^n; \mathbb{Z}_2)$ are $[a + \partial b]$ and $[\partial b]$. The chain map $\bar{f}_\sharp$ sends singular 0-simplices to singular 0-simplices, so $\bar{f}_*$ is the identity. In particular, it's an isomorphism.

Suppose $\bar{f}_*$ is an isomorphism from $H_{i-1}(P^n)$ to itself. Then if we invert the isomorphism on the bottom of the left commutative square, we get that $\bar{f}_*$ is also an isomorphism from $H_i(P^n)$ to itself. This applies all the way to $i = n$, so if we invert the isomorphism on the top of right commutative square, we get that f_* is an isomorphism from $H_n(S^n)$ to itself.

We now want to translate the degree of f to $\mathbb{Z}_2$ coefficients. If $\phi : \mathbb{Z} \to \mathbb{Z}_2$ is a homomorphism, it induces a chain map $\phi_\sharp : C_i(S^n; \mathbb{Z}) \to C_i(S^n; \mathbb{Z}_2)$, and therefore a homomorphism $\phi_* : H_i(S^n; \mathbb{Z}) \to H_i(S^n; \mathbb{Z}_2)$. This homomorphism commutes with the homomorphisms induced by maps $S^n \to S^n$, and we haven't built the framework in this paper to show this, but ϕ_* also gives a commutative diagram mapping the long exact sequences derived from Theorem 2.26. In particular, we have this commutative diagram: (See the proof of Lemma 3.4)

$$\begin{tikzcd}
\tilde{H}_i(S^i;\mathbb{Z}) \arrow[r, "\approx"] \arrow[d, "\phi_*"] & \tilde{H}_{i-1}(S^{i-1};\mathbb{Z}) \arrow[d, "\phi_*"] \\
\tilde{H}_i(S^i;\mathbb{Z}_2) \arrow[r, "\approx"] & \tilde{H}_{i-1}(S^{i-1};\mathbb{Z}_2)
\end{tikzcd}$$

Using these facts, and inducting from the case $n = 0$ where

$$\tilde{H}_0(S^0;G) = \{g\sigma_1 - g\sigma_2 \ : \ g \in G,\ \sigma_1, \sigma_2 : \Delta^0 \to S^0\},$$

we find that this diagram commutes:

$$\begin{tikzcd}
\mathbb{Z} \arrow[r, "\approx"] \arrow[d, "\phi"] & \tilde{H}_n(S^n;\mathbb{Z}) \arrow[r, "f_*"] \arrow[d, "\phi_*"] & \tilde{H}_n(S^n;\mathbb{Z}) \arrow[r, "\approx"] \arrow[d, "\phi_*"] & \mathbb{Z} \arrow[d, "\phi"] \\
\mathbb{Z}_2 \arrow[r, "\approx"] & \tilde{H}_n(S^n;\mathbb{Z}_2) \arrow[r, "f_*"] & \tilde{H}_n(S^n;\mathbb{Z}_2) \arrow[r, "\approx"] & \mathbb{Z}_2
\end{tikzcd}$$

Going along the bottom, where f_* is an isomorphism, sends $1 \in \mathbb{Z}$ to $1 \in \mathbb{Z}_2$. But going across the top sends $1 \in \mathbb{Z}$ to $\deg f$ modulo 2, so f has odd degree. $\square$

Now, for the proof of the Borsuk-Ulam Theorem:

Proof. Suppose that the theorem is false. That is, suppose there is a function $f : S^n \to \mathbb{R}^n$ where $f(\boldsymbol{x}) \neq f(-\boldsymbol{x})$ for all $\boldsymbol{x}$. Then we can construct a function $g : S^n \to S^{n-1}$ by

$$g(\boldsymbol{x}) \coloneqq \frac{f(\boldsymbol{x}) - f(-\boldsymbol{x})}{\|f(\boldsymbol{x}) - f(-\boldsymbol{x})\|}$$

Let h be the restriction of g to the equator, S^{n-1}.

One can check that h is odd, so by Lemma 3.7, h has odd degree.

However, h is also homotopic to a constant map. To see this, note that the closed upper hemisphere of S^n (which can be thought of as the set of points in S^n with nonnegative last coordinate) is homeomorphic to B^n, and hence there is a homotopy $F(x,t)$ between the identity on the upper hemisphere and a constant map. We can construct a homotopy between h and a constant map with

$$G(x,t) \coloneqq h(F(x,t))$$

Homotopic maps induce the same homomorphism on homology groups, and a constant map has degree 0, so h also has degree 0. But this contradicts our previous finding that h has odd degree. Hence, the function f cannot exist. $\square$

4. Tucker's Lemma and the Hex Theorem

We'll explain each result, prove its equivalence to the corresponding continuous result, and prove each result independently. Then, we'll show how Tucker's Lemma implies the Hex Theorem, completing the diagram which connects the four theorems in this paper.

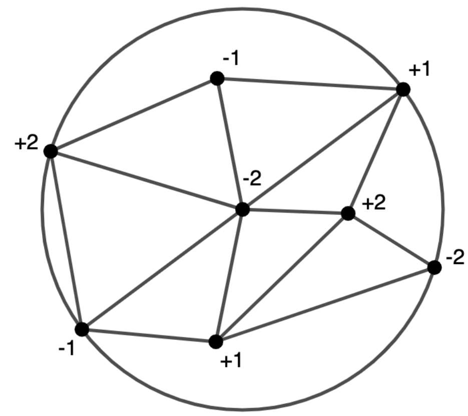

FIGURE 2. An example of a Tucker triangulation and labeling

4.1. **Tucker's Lemma.** Tucker's Lemma is a surprising result about triangulations, which uses information about the triangulation at the boundary to deduce a property that manifests in the interior. First, define the following kind of triangulation:

Definition 4.1. Suppose T is a triangulation of B^n (or, where appropriate, a homeomorphic subset of $\mathbb{R}^n$), together with a labeling $\lambda : V(T) \to \{\pm 1, \ldots, \pm n\}$. Then the pair (T, λ) is a **Tucker pair** if T and λ are both antipodal on the boundary. That is:

(1) For all $\sigma \in T$ such that $\sigma \subset \partial B^n$, we have $-\sigma \in T$.
(2) For all vertices $\boldsymbol{v} \in V(T)$ such that $\boldsymbol{v} \in \partial B^n$, we have $\lambda(-\boldsymbol{v}) = -\lambda(\boldsymbol{v})$.

We call such a T a **Tucker triangulation**, and call such a λ a **Tucker labeling**.

Note that this definition is only concerned with the behavior of T and λ on the boundary. This makes Tucker's Lemma quite remarkable:

Theorem 4.2. *Suppose (T, λ) is a Tucker pair. Then there is a 1-simplex (edge) in the triangulation which is* ***complementary*** *(its vertices' labels are negatives of each other).*

Tucker's Lemma, like the Borsuk-Ulam Theorem, has an equivalent formulation negating the existence of a kind of embedding $B^n \to S^{n-1}$. To state it, let's first define a specific family of simplicial complexes:

Definition 4.3. For each n, let $\Diamond^{n-1}$ denote the simplicial complex with vertex set $\{\pm 1, \ldots, \pm n\}$, and with $\sigma \subseteq V(\Diamond^{n-1})$ a simplex when σ does not contain both i and $-i$ for any i. Geometrically, we associate the vertex $+i$ with $\boldsymbol{e}_i$, and the vertex $-i$ with $-\boldsymbol{e}_i$, so $\left\|\Diamond^{n-1}\right\| = \{\boldsymbol{x} \in \mathbb{R}^n \ : \sum_i |x_i| = 1\}$.

Then we can rephrase Tucker's Lemma as:

Theorem 4.4. *Suppose T is a Tucker triangulation. Then there is no simplicial map from T to $\Diamond^{n-1}$ which is odd on the boundary.*

That these two formulations are equivalent comes from identifying λ in the first formulation with any attempted antipodal simplicial map in the second formulation.

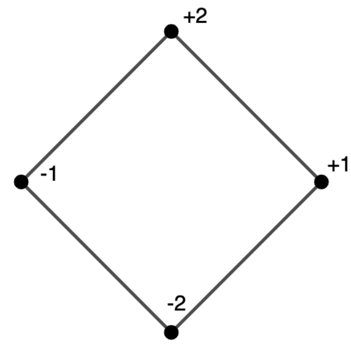

FIGURE 3. The inspiration for the name – $\diamond^1$

Proving Tucker's Lemma from the Borsuk-Ulam Theorem is straightforward, using the nonexistence versions of both (Theorems 3.2 and 4.4):

Proof of Theorem 4.4 from Theorem 3.2. Suppose we have a Tucker triangulation T and a simplicial map $\lambda : V(T) \to V(\diamond^{n-1})$ which is odd on the boundary. Then the affine extension of λ can be composed with central projection to get a continuous function $B^n \to S^{n-1}$ which is odd on the boundary. Our desired result is the contrapositive of what we've just shown. □

The other direction will be somewhat constructive, using Tucker's Lemma to find points which get mapped arbitrarily close to their antipodal pair. Compactness gives us the full result.

Let's quickly introduce some notation:

Definition 4.5. The ℓ_∞ **norm**, or $\|\cdot\|_\infty$, is defined by

$$\|(x_1, x_2, \dots, x_n)\|_\infty \coloneqq \max\{|x_i|\}$$

This norm produces much of the same topological structures on $\mathbb{R}^n$ as the standard norm. In particular, they produce the same compact sets, and a function is uniformly continuous under one norm if and only if it's uniformly continuous under the other.

Proof of Theorem 3.1 from Theorem 4.2. Let $g : S^n \to \mathbb{R}^n$ be a continuous function. If we let $f(\boldsymbol{x}) \coloneqq g(\boldsymbol{x}) - g(-\boldsymbol{x})$, then we wish to show that $f(\boldsymbol{x}) = 0$ for some $\boldsymbol{x}$. Note that f is odd and continuous.

In this proof, we will triangulate the closed upper hemisphere of S^n. Although this isn't B^n, it is homeomorphic to it, and its boundary has antipodal symmetry, so Tucker's Lemma still applies.

Let $\varepsilon > 0$. Because S^n is compact, f must be uniformly continuous, so there is a $\delta > 0$ such that $\|f(\boldsymbol{x}) - f(\boldsymbol{y})\|_\infty < \varepsilon$ whenever $\|\boldsymbol{x} - \boldsymbol{y}\|_\infty < \delta$. We let T be a Tucker triangulation of the upper hemisphere where every simplex in T has diameter at most δ under the ℓ^∞ norm.

We now define the Tucker labeling λ. First, let

$$k(\boldsymbol{v}) \coloneqq \min\{i \ : \ |f(\boldsymbol{v})_i| = \|f(v)\|_\infty\}$$

for every vertex $\boldsymbol{v}$ in T. We define $\lambda : V(T) \to \{\pm 1, \ldots, \pm n\}$ by

$$\lambda(\boldsymbol{v}) \coloneqq \begin{cases} +k(\boldsymbol{v}), & f(\boldsymbol{v})_{k(\boldsymbol{v})} > 0 \\ -k(\boldsymbol{v}), & f(\boldsymbol{v})_{k(\boldsymbol{v})} < 0 \end{cases}$$

This labeling represents the direction in which $\boldsymbol{v}$ gets sent the farthest from the origin.

Since f is odd, λ is odd on the boundary. The pair (T, λ) is therefore a Tucker pair, and contains a complementary edge $\boldsymbol{v}\boldsymbol{v}'$. WLOG, suppose $i \coloneqq \lambda(\boldsymbol{v}) > 0$. Then, by our construction of T, we get:

$$\begin{aligned} \|f(\boldsymbol{v}) - f(\boldsymbol{v}')\|_\infty &< \varepsilon \\ |f(\boldsymbol{v})_i - f(\boldsymbol{v}')_i| &< \varepsilon \\ f(\boldsymbol{v})_i - f(\boldsymbol{v}')_i &< \varepsilon \\ f(\boldsymbol{v})_i &< \varepsilon \\ \|f(\boldsymbol{v})\|_\infty &< \varepsilon \end{aligned}$$

So, for any $\varepsilon > 0$, there is some $\boldsymbol{x} \in S^n$ such that $\|f(\boldsymbol{x})\|_\infty < \varepsilon$. We can construct a sequence $\{\boldsymbol{x}_n\}$ of these points, so that $\lim_{n\to\infty} f(\boldsymbol{x}_n) = 0$. Because S^n is compact, this sequence has a convergent subsequence. We might as well just throw out the rest of the sequence, and let $\{\boldsymbol{x}_n\}$ be the convergent subsequence. Let $\boldsymbol{x}^* = \lim_{n\to\infty} \boldsymbol{x}_n$. Then $f(\boldsymbol{x}^*) = 0$ by continuity. $\square$

The proof of Tucker's Lemma we'll present is constructive, so combining these two proofs will give a method for finding antipodal points mapped arbitrarily close to each other. (Performing such an algorithm usually doesn't require the entire Tucker triangulation to be found, so it's computationally feasible.)

4.2. **Proof of Tucker's Lemma.** The proofs of Theorems 4.2 and 4.4 are all topological. In order to prove Tucker's Lemma with combinatorial techniques, we have to weaken it, and place further conditions on a Tucker triangulation. These conditions will still allow the simplices in a Tucker triangulation to have arbitrarily small diameter, so the proof that Tucker's Lemma implies the Borsuk-Ulam Theorem still holds. This means that this "weaker" version of Tucker's Lemma is equivalent to the complete version.

Let $\hat{B}^n$ be the unit ball in $\mathbb{R}^n$ under the ℓ_1 norm. That is,

$$\hat{B}^n \coloneqq \{(x_1, \ldots, x_n) \in \mathbb{R}^n \ : \ |x_1| + \cdots + |x_n| \le 1\}$$

Our condition on Tucker's Lemma will involve this natural triangulation of $\hat{B}^n$:

Definition 4.6. $\diamondplus^n$ is a simplicial complex with vertex set

$$V(\diamondplus^n) \coloneqq \{\mathbf{0}, \pm 1, \pm 2, \ldots, \pm n\}$$

and $\sigma \subseteq V(\diamondplus^n)$ a simplex if and only if it doesn't contain both $+i$ and $-i$ for any $i \in \{1, \ldots, n\}$.

Every simplex in $\diamondplus^n$ is either a simplex in $\diamond^{n-1}$ or the union of a simplex in $\diamond^{n-1}$ with the set $\{\mathbf{0}\}$. See figure 4 for what $\diamondplus^2$ looks like.

We can embed $\diamondplus^n$ in $\mathbb{R}^n$ by putting $\mathbf{0}$ at the origin, $+i$ at the unit vector $\boldsymbol{e}_i$, and $-i$ at the unit vector $-\boldsymbol{e}_i$. In this case, the sign of each coordinate is constant within the relative interior of each simplex.

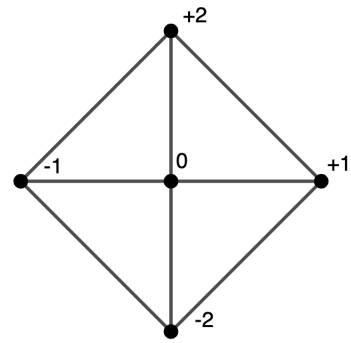

FIGURE 4.

Definition 4.7. A **special Tucker triangulation** is a triangulation of $\hat{B}^n$ which is antipodal on the boundary of $\hat{B}^n$ and refines $\diamond^n$. That is, every simplex in a special Tucker triangulation is contained in a simplex of $\diamond^n$.

An equivalent way of stating the second part of the definition is that the sign of each coordinate is constant within the relative interior of each simplex. Note that $\hat{B}^n$ is homeomorphic to B^n, so such a triangulation is also a normal Tucker triangulation.

We will prove that Tucker's Lemma is true for all Tucker pairs (T, λ) where T is special. The proof is taken from [1], with the proper adjustments made to make it constructive.

Theorem 4.8. *If (T, λ) is a Tucker pair where T is a special Tucker triangulation, then there is a 1-simplex in T which is complementary.*

Proof. First, some notation. For a simplex σ, let

$$\lambda(\sigma) := \{\lambda(\boldsymbol{v}) \ : \ \boldsymbol{v} \in V(\sigma)\}$$

and let

$$S(\sigma) := \{i \ : \ x_i > 0\} \cup \{-i \ : \ x_i < 0\}$$

where $(x_1, \ldots, x_n)$ is a point in the relative interior of σ. This is well-defined because T is a special triangulation. Call a simplex σ "happy" if $S(\sigma) \subseteq \lambda(\sigma)$.

If σ is a happy simplex, let $k := |S(\sigma)|$. We have that σ is contained in the span of the vectors $\boldsymbol{e}_{|i|}$, where $i \in S(\sigma)$, so σ is contained in a k-dimensional subspace of $\mathbb{R}^n$. Hence we must have $\dim \sigma \leq k$. But at least k vertex labels are required for σ to be happy, so $\dim \sigma \geq k - 1$. If $\dim \sigma = k - 1$, so that σ needs every vertex label to be happy, then we'll call σ "tight". If $\dim \sigma = k$, then we'll call σ "loose".

We quickly note three properties of happy simplices. First, a happy simplex σ on the boundary must be tight, because then $|S(\sigma)|$ is one more than the dimension of the simplex in $\diamond^n$ whose relative interior contains σ's. Second, the simplex $\{\mathbf{0}\}$ is happy and loose, because $S(\{\mathbf{0}\}) = \emptyset$. And third, the empty simplex is happy and tight.

We define a graph of the happy simplices of T, where two simplices σ and τ are connected if:

- σ is on the boundary of $\hat{B}^n$ and τ is its antipodal partner ($\tau = -\sigma$), or

- σ is a facet of τ with $\lambda(\sigma) = S(\tau)$

We want to show that every happy simplex has degree at most 2, the empty simplex has degree 1, and any other simplex with degree 1 contains a complementary edge. Then by Theorem 2.1, there is a path starting from the empty simplex which ends in a simplex containing a complementary edge. (The empty simplex's neighbor is always $\{\mathbf{0}\}$, so you may as well consider the path to start from $\{\mathbf{0}\}$ for computations.)

Let σ be happy simplex. We distinguish two cases:

(1) σ is tight. Then it cannot have a facet τ with $\lambda(\tau) = S(\sigma)$, because then τ would have at least the same dimension as σ. If σ is on the boundary, then it's neighbors with its antipodal pair, but is only the facet of one happy simplex. If it's not on the boundary, then it is the facet of two happy simplices. These happy simplices τ are in a k-dimensional linear subspace of $\mathbb{R}^n$ with σ, so $S(\tau) = S(\sigma) = \lambda(\sigma)$. Hence, τ is a neighbor of σ. One exception to this reasoning is the empty simplex. Its k-dimensional subspace is just the origin, so it only has one neighbor, the simplex $\{\mathbf{0}\}$. In conclusion, every tight simplex has exactly two neighbors – except for the empty simplex, which only has one.

(2) σ is loose. Then it either has a repeated vertex label, or a vertex label i which doesn't appear in $S(\sigma)$. If the former, then σ can't be the facet of a happy simplex (because $|\lambda(\sigma)| = k$), but is neighbors with its two happy facets. If the latter, then σ is neighbors with its facet which doesn't contain i, and no other facets. Moreover, if $-i$ isn't also in $\lambda(\sigma)$, then σ is a facet of a simplex τ with $S(\tau) = S(\sigma) \cup \{i\}$. This simplex can be reached from σ by moving along the i-th coordinate axis (in the positive direction if $i > 0$, and in the negative direction if $i < 0$). If $-i$ is also in $\lambda(\sigma)$, then it is in $S(\sigma)$, and so no such simplex τ exists, because T is special. Note that the former case of a loose simplex cannot contain a complementary edge, because $S(\sigma)$ cannot contain i and $-i$. In conclusion, a loose simplex has exactly two neighbors – unless it contains a complementary edge, in which case it only has one.

As we said earlier, this guarantees a path from $\{\mathbf{0}\}$ to a simplex with a complementary edge. This is the only path from $\{\mathbf{0}\}$ (other than the path to the empty simplex), so you can use the graph's definition to find the path algorithmically. $\square$

4.3. **The Hex Theorem.** This theorem gets its name from a game called "Hex". It's a two-player game with a board made of hexagonal tiles, where players alternate placing their color on a hexagon, until one player forms a connected path of hexagons from one side of the board to another. Which pair of sides each player has to connect is dependent on which color they are. For the purpose of mathematical analysis, the game is usually represented as a game where players color a graph like in Figure 5.

The Hex Theorem traditionally says that at least one player must always win. The game of Hex also generalizes to n dimensions and players, and the Hex Theorem can be shown to be true for these specific graphs (see [3]). In this paper, we will generalize a bit further, by defining a Hex board more broadly than the definition usually given (a graph where $\boldsymbol{x}$ and $\boldsymbol{y}$ are connected if $\|\boldsymbol{x} - \boldsymbol{y}\|_\infty = 1$). This generalization is inspired by a previous paper from the UChicago REU [4].

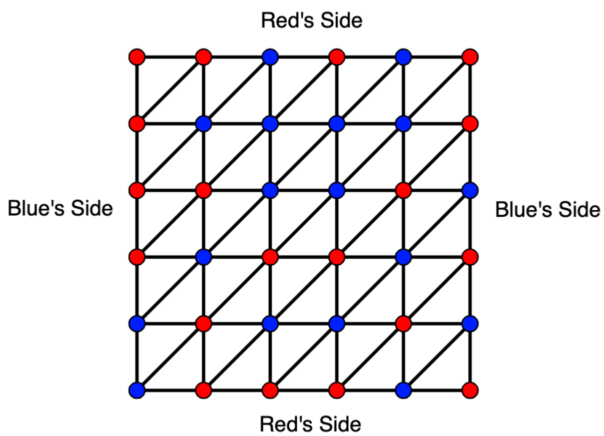

FIGURE 5. Traditional Hex board as a graph

First, here's how the Hex Theorem looks in its traditional form:

Theorem 4.9. *If every vertex on a two-player Hex board is colored red or blue, then there is a path of red vertices from the top of the board to the bottom, or there is a path of blue vertices from the left side of the board to the right side.*

The intuition for the generalization is taken from how this and the n-dimensional Hex Theorem are proven, so it may be useful to keep the normal Hex board in mind, or consult [3] or [4] for a more traditional treatment.

First, we define a Hex game:

Definition 4.10. An n-dimensional **Hex game** Γ is a pair (T, λ) where T is a simplicial complex with $\|T\| = I^n$ and λ is a coloring function sending $V(T)$ to $\{1, \ldots, n\}$. T is called a **Hex board**.

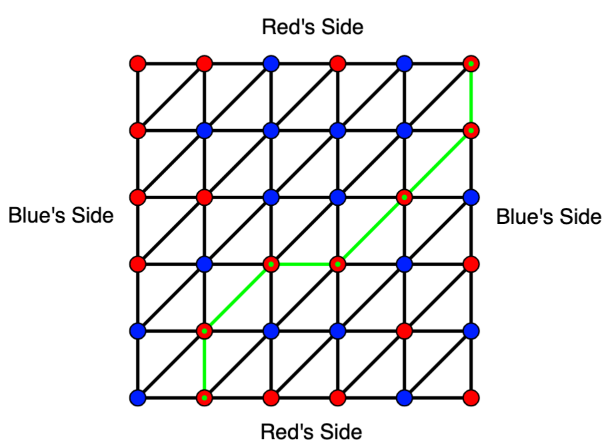

FIGURE 6. Hex board in Figure 5, with winning path highlighted

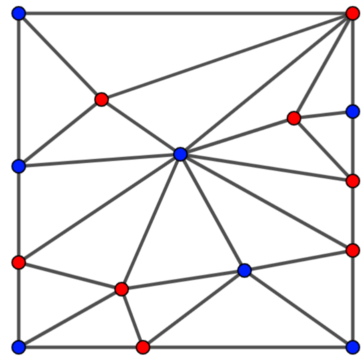

Figure 7. Example of a general 2-dimensional Hex board

A central part of the Hex Theorem is the idea of connectedness:

Definition 4.11. Two vertices $\boldsymbol{a}$, $\boldsymbol{b}$ in a Hex game (T, λ) are **connected** if there is a sequence of 1-simplices in T of the form

$$[\boldsymbol{a}, \boldsymbol{v}_1], [\boldsymbol{v}_1, \boldsymbol{v}_2], \ldots, [\boldsymbol{v}_{k-1}, \boldsymbol{v}_k], [\boldsymbol{v}_k, \boldsymbol{b}]$$

We say these two vertices are ***i*-connected** if $\lambda(\boldsymbol{a}) = \lambda(\boldsymbol{b}) = i$ and $\lambda(\boldsymbol{v}_j) = i$ for all j between 1 and k. We call such a sequence an ***i*-path**.

The last definition required is the sides to be connected to win:

Definition 4.12.

$$F(\{+i\}) \coloneqq \{\boldsymbol{x} \in V(T) \ : \ x_i = 0\}$$
$$F(\{-i\}) \coloneqq \{\boldsymbol{x} \in V(T) \ : \ x_i = 1\}$$

And in general, $F : 2^{\{\pm 1, \ldots, \pm n\}} \to 2^{V(T)}$ is defined by

$$F(\{i_1, i_2, \ldots, i_k\}) \coloneqq \bigcap_{m=1}^{k} F(\{i_m\})$$

At last, here is the Hex Theorem:

Theorem 4.13. *In every Hex game, for some color i, there are $\boldsymbol{a} \in F(\{+i\})$ and $\boldsymbol{b} \in F(\{-i\})$ which are i-connected.*

The proof of Brouwer's Fixed Point Theorem from the Hex Theorem is akin to the proof of the Borsuk-Ulam Theorem from Tucker's Lemma. For the equivalence, we will use the fact that I^n and B^n are homeomorphic to make Brouwer's Fixed Point Theorem concern functions $f : I^n \to I^n$.

Proof of Brouwer's Fixed Point Theorem from the Hex Theorem. Let f be a continuous function from I^n into itself, and let $\varepsilon > 0$.

Because I^n is compact, f is uniformly continuous, so there is $\delta > 0$ smaller than ε such that $\|f(\boldsymbol{x}) - f(\boldsymbol{y})\|_\infty < \varepsilon$ whenever $\|\boldsymbol{x} - \boldsymbol{y}\|_\infty < \delta$. Let T be a Hex board where the diameter of every simplex in T is less than δ, under the ℓ_∞ norm.

For any $\boldsymbol{x} \in I^n$, if i is the smallest index such that $\|f(\boldsymbol{x}) - \boldsymbol{x}\|_\infty = |f_i(\boldsymbol{x}) - x_i|$, we'll say that $\boldsymbol{x}$ is "moved in the direction i". Define the Hex coloring $\lambda(\boldsymbol{x})$ to be the direction $\boldsymbol{x}$ is moved in.

By the Hex Theorem, for some i, there is an i-path connecting $F(\{+i\})$ and $F(\{-i\})$. For a vertex $\boldsymbol{x} \in F(\{+i\})$, we must have $f_i(\boldsymbol{x}) - x_i \geq 0$, while for a vertex $\boldsymbol{x} \in F(\{-i\})$, we have $f_i(\boldsymbol{x}) - x_i \leq 0$. One of these two statements is true for every vertex on the i-path. Then somewhere on the path is an edge $[\boldsymbol{x}, \boldsymbol{y}]$ with $f_i(\boldsymbol{x}) - x_i \geq 0$ and $f_i(\boldsymbol{y}) - y_i \leq 0$. It cannot be the case that both $\|f(\boldsymbol{x}) - \boldsymbol{x}\|_\infty \geq \varepsilon$ and $\|f(\boldsymbol{y}) - \boldsymbol{y}\|_\infty \geq \varepsilon$, because then

$$\begin{aligned}
2\varepsilon &\leq \|f(\boldsymbol{x}) - \boldsymbol{x}\|_\infty + \|f(\boldsymbol{y}) - \boldsymbol{y}\|_\infty \\
&= (f_i(\boldsymbol{x}) - x_i) + (y_i - f_i(\boldsymbol{y})) \\
&= (f_i(\boldsymbol{x}) - f_i(\boldsymbol{y})) + (y_i - x_i) \\
&\leq \|f(\boldsymbol{x}) - f(\boldsymbol{y})\|_\infty + \|\boldsymbol{y} - \boldsymbol{x}\|_\infty \\
&< \varepsilon + \delta \\
&< 2\varepsilon
\end{aligned}$$

So, for every $\varepsilon > 0$, we can use the Hex Theorem to find a point $\boldsymbol{x}$ where $\|f(\boldsymbol{x}) - \boldsymbol{x}\|_\infty < \varepsilon$. Therefore, because I^n is compact, it contains a convergent sequence $\{\boldsymbol{x}_n\}$ where $\lim_{n\to\infty} f(\boldsymbol{x}_n) = \lim_{n\to\infty} \boldsymbol{x}_n$. By continuity, if $\boldsymbol{x}^*$ is the limit of the sequence, then $f(\boldsymbol{x}^*) = \boldsymbol{x}^*$. So f has a fixed point. □

Proof of the Hex Theorem from Brouwer's Fixed Point Theorem. Suppose there is a Hex game (T, λ) such that for each i, there is no i-path connecting $F(\{+i\})$ and $F(\{-i\})$. We'll first construct a function from $V(T)$ to I^n which preserves simplices.

Define $f : V(T) \to I^n$ as follows:

(1) If $\lambda(\boldsymbol{v}) = 1$, then send $\boldsymbol{v}$ into the relative interior of a simplex of which it is a face, moving it along $\boldsymbol{e}_1$ if there is a 1-path connecting $\boldsymbol{v}$ to $F(\{+1\})$, and along $-\boldsymbol{e}_1$ if there isn't. By our assumptions about the Hex game, this produces a geometric realization of the same triangulation T. The amount moved doesn't have to be the same for all such $\boldsymbol{v}$, but it can be.

(2) If we've defined f for vertices colored $i-1$, then for a vertex $\boldsymbol{v}$ colored i, send $\boldsymbol{v}$ into the relative interior of a simplex of which it is a face, moving it along $\boldsymbol{e}_i$ if there is an i-path connecting $\boldsymbol{v}$ to $F(\{+i\})$, and along $-\boldsymbol{e}_i$ if there isn't.

The function f is a simplicial mapping, so we can consider its affine extension $\|f\| : I^n \to I^n$. We show a lemma now.

Lemma 4.14. *If $f([\boldsymbol{x}_0, \ldots, \boldsymbol{x}_n]) = [\boldsymbol{y}_0, \ldots, \boldsymbol{y}_n]$, then $\|f\|$ has a fixed point in this simplex if and only if $\mathbf{0} \in \operatorname{conv}\{\boldsymbol{y}_0 - \boldsymbol{x}_0, \ldots, \boldsymbol{y}_n - \boldsymbol{x}_n\}$.*

Proof. A point in the domain simplex can be written as $\boldsymbol{x} = \sum_{i=0}^n \alpha_i \boldsymbol{x}_i$, where the coefficients are nonnegative and sum to 1. By the definition of $\|f\|$, we have:

$$\begin{aligned}
\|f\|(\boldsymbol{x}) - \boldsymbol{x} &= \sum_{i=0}^n \alpha_i \boldsymbol{y}_i - \sum_{i=0}^n \alpha_i \boldsymbol{x}_i \\
&= \sum_{i=0}^n \alpha_i (\boldsymbol{y}_i - \boldsymbol{x}_i)
\end{aligned}$$

So $\|f\|(\boldsymbol{x}) - \boldsymbol{x}$ is a point in $\operatorname{conv}\{\boldsymbol{y}_0 - \boldsymbol{x}_0, \ldots, \boldsymbol{y}_n - \boldsymbol{x}_n\}$, and every point in that convex hull can be expressed as $\|f\|(\boldsymbol{x}) - \boldsymbol{x}$ for some $\boldsymbol{x} \in \operatorname{conv}\{\boldsymbol{x}_0, \ldots, \boldsymbol{x}_n\}$. Hence, $\|f\|(\boldsymbol{x}) = \boldsymbol{x}$ for such an $\boldsymbol{x}$ if and only if $\mathbf{0} \in \operatorname{conv}\{\boldsymbol{y}_0 - \boldsymbol{x}_0, \ldots, \boldsymbol{y}_n - \boldsymbol{x}_n\}$. □

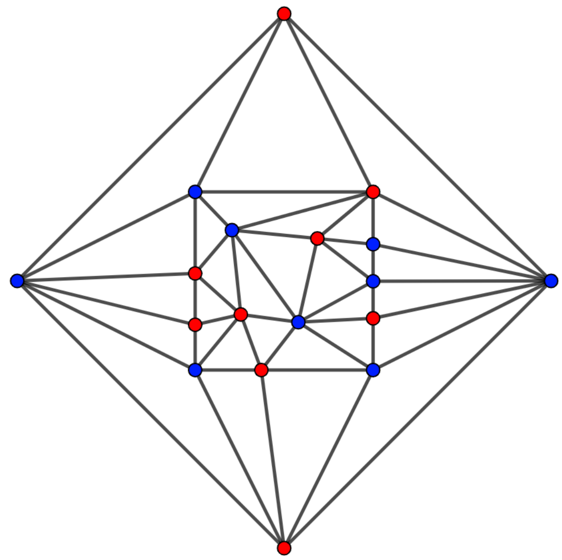

FIGURE 8. Example of a 2-dimensional augmented Hex board

In the way we've defined f, the vertices of the simplices in T get moved along standard unit vectors, and within one simplex, no two vertices are moved in opposite directions. Our lemma then implies that $\|f\|$ doesn't have a fixed point in any simplex in T. That is, $\|f\|$ is a continuous function from I^n into itself with no fixed point. $\square$

4.4. **Proof of the Hex Theorem.** Our constructive proof will use a modified Hex game, so that the boundary of the board always looks the same. This will allow us to perform a similar sort of constructive proof as we did for Tucker's Lemma.

Definition 4.15. If $\Gamma = (T, \lambda)$ is a Hex game, then the **augmented Hex game** $\hat{\Gamma} = (\hat{T}, \lambda)$ is defined as follows:

- $V(\hat{T}) \coloneqq V(T) \cup \{\pm 1, \ldots, \pm n\}$, where these are assumed to be new vertices.
- Every simplex in T is a simplex in $\hat{T}$
- If $\sigma \subseteq F(\tau)$ is a simplex in T, where τ doesn't contain both i and $-i$, then $\sigma \cup \tau$ is a simplex in $\hat{T}$.
- $\lambda(i) \coloneqq |i|$, where $i \in \{\pm 1, \ldots, \pm n\}$.

Note that the added simplices include every subset of $\{\pm 1, \ldots, \pm n\}$ which doesn't include a number and its negative, because $\sigma \subseteq F(\tau)$ can be empty. The augmented Hex game is then a kind of suspension of the Hex game inside $\diamondsuit^{n-1}$.

Proof of the Hex Theorem. Call a simplex in $\hat{\Gamma}$ "completely colored" if λ sends its vertices onto all of $\{1, \ldots, n\}$. We define a graph with the completely colored simplices in $\hat{\Gamma}$ as the vertices, where two simplices are connected if one is a completely colored facet of the other.

If σ is completely colored, then its dimension is n or $n-1$. If the former is the case, then σ has exactly two vertices which are the same color, and thus has two completely colored facets. If the latter is the case, then σ is a completely colored facet of one or two completely colored simplices.

So every vertex in the graph has degree at most 2, and by Theorem 2.1 the graph is composed of simple paths, simple cycles, and isolated vertices. The only completely colored simplices which are facets of just one completely colored simplex are those on the boundary, $\diamond^{n-1}$, and all the $(n-1)$-dimensional simplices on the boundary are completely colored, so they are the only vertices in the graph with degree 1. Therefore, there is a path from the simplex $[1, \ldots, n]$ to some other completely colored simplex in $\diamond^{n-1}$. Each pass in the graph from one completely colored facet to the next connects the vertices of the same color. So a path from $[1, \ldots, n]$ to another face of $\diamond^{n-1}$ i-connects the vertex $+i$ with either $+i$ or $-i$. If the face is different from $[1, \ldots, n]$, then it must have a negative vertex, $-i$. So our path i-connects $+i$ and $-i$.

The only vertices in $\hat{\Gamma}$ colored i, other than the vertices $+i$ and $-i$, are those colored i in the original game Γ. So the only i-colored vertices the vertex $+i$ shares an edge with are in $F(\{+i\})$, and the only i-colored vertices the vertex $-i$ shares an edge with are in $F(\{-i\})$. Therefore, the i-colored path connecting $+i$ and $-i$ also i-connects a vertex in $F(\{+i\})$ and a vertex in $F(\{-i\})$. $\square$

4.5. **Proof that Tucker's Lemma Implies the Hex Theorem.** Consider the augmented Hex Board $\hat{\Gamma}$. This is a triangulation of $\hat{B}^n$. If we label the origin $\mathbf{0}$, then we can geometrically realize $\hat{T}$ as a special Tucker triangulation by putting $+i$ at $-2\boldsymbol{e}_i$ and $-i$ at $2\boldsymbol{e}_i$. (Specifying it's a special triangulation means we can use the constructive proof when we apply Tucker's Lemma.)

We now define a Tucker labeling $\hat{\lambda}$ on this triangulation. Suppose $\lambda(\boldsymbol{v}) = i$. Then we let

$$\hat{\lambda}(\boldsymbol{v}) \coloneqq \begin{cases} -i, & \boldsymbol{v} = -i, \text{ or } \boldsymbol{v} \text{ and } +i \text{ aren't } i\text{-connected} \\ +i, & \text{otherwise} \end{cases}$$

Then $(\hat{T}, \hat{\lambda})$ is a Tucker pair, so $\hat{T}$ contains a complementary edge. This edge cannnot be in the original triangulation T, nor can it involve the vertex $+i$, as those would contradict the definition of $\hat{\lambda}$. So, the vertex in the edge with label $-i$ must be $-i$, and this implies that $-i$ and $+i$ are i-connected.

Acknowledgments

I'd like to give my deepest thanks to my mentor, Peng Hui How, who introduced me to the topic of this paper, and helped me get a working understanding of basic algebraic topology. I would also like to thank Professor Peter May for organizing a hybrid REU, Professor Daniil Rudneko for running the Apprentice program, and all the REU speakers for their engaging talks.

References

[1] J. Matoušek. *Using the Borsuk-Ulam Theorem.* Springer-Verlag Berlin Heidelberg, 2003.
[2] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2001. Available at `http://pi.math.cornell.edu/~hatcher/AT/AT.pdf`.
[3] D. Gale. *The Game of Hex and the Brouwer Fixed-Point Theorem.* The American Mathematical Monthly, Vol. 86 No. 10, pp. 818-827.

[4] M. Shachner. *The Game of Hex: A Study in Graph Theory and Algebraic Topology*. Online; Published August 29, 2020. Available at http://math.uchicago.edu/~may/REU2019/REUPapers/Schachner.pdf.

[5] V. Guillemin and A. Pollack. *Differential Topology*. Prentice-Hall, Inc., 1974.

[6] http://math.mit.edu/~fox/MAT307-lecture03.pdf.

[7] https://web.mit.edu/sp.268/www/hex-notes.pdf.