Topological methods in combinatorics and game theory

Gabriel Coutinho

February 26, 2021

These are the course notes of a (under)grad course being offered at UFMG in 2020.2.

In 2013 I took an inspiring course on Topological Methods with Prof. Penny Haxell in the C&O Department in Waterloo. A significant amount of this course will follow hers almost to the letter. Several sections of these course notes are copied from her class notes (with her permission, of course). These will be marked with a star on the title.

Further sources for this course notes are

- Jiří Matoušek's great book "Using the Borsuk-Ulam Theorem".
- Laszlo Lovász's notes on Topological Methods in Combinatorics.
- Jesus Deloera's notes on Topological Methods in Game Theory.

This course will not exactly assume prior knowledge in topology, graph theory or game theory. However, whenever I introduce a concept that is new to you, I recommend you research it on your own (wikipedia page should suffice, most of the time). This is basically to ensure you understand the motivation behind what we are doing.

Contents

1 Motivation

I start this section by presenting two theorems.

Theorem (Brouwer 1910). *Let C be a nonempty, compact, convex subset of $\mathbb{R}^n$, and $f : C \to C$ a continuous function. Then f admits a fixed point, that is, there is $x \in C$ so that $f(x) = x$.*

Let S_r denote the sphere (according to the Euclidean distance) in $\mathbb{R}^{r+1}$. Thus S_1 is a circle in $\mathbb{R}^2$, for example.

Theorem (Borsuk-Ulam 1933). *For every continuous map $f : S_r \to \mathbb{R}^r$, there exists $x \in S_r$ so that $f(x) = f(-x)$.*

Both these theorems have combinatorial counterparts — their statements however will come later. Sperner's Lemma is that of Brouwer's Theorem, and if fact it is called a lemma because it can be used to show the latter. Tucker's Lemma is the combinatorial counterpart of Borsuk-Ulam's, and again, it can be used as lemma in its proof.

More interestingly, Sperner's Lemma and Tucker's Lemma have found numerous applications in combinatorics, and one of the goals of this course is to explore these applications. One of the results we will prove using Sperner's Lemma is the following.

Let G be a graph, and let $\mathcal{P} = V_1 \cup ... \cup V_m = V(G)$ be a partition of $V(G)$. Then an *independent transversal* of G with respect to $\mathcal{P}$ is a set $S = \{v_1, ..., v_m\} \subseteq V(G)$ such that

- S is independent (that is, the subgraph $G[S]$ induced by the vertices of S contains no edges);
- $v_i \subseteq V_i$, for all i.

The problem of finding independent transversals is not easy in general — in fact, it is easy[1] to show a reduction of the Exact Cover problem (its decision version is one of Karp's original 21 NP-complete problems) to that of finding an independent transversal.

Let $\Delta = \Delta(G)$ denote the maximum degree of G.

Theorem. *Let G be a graph and $\mathcal{P}$ a vertex partition of G. Suppose $|V_i| \geq 2\Delta$ for all i. Then G has an independent transversal with respect to $\mathcal{P}$.*

To understand how a theorem that relates to topology can be used to show the purely combinatorial result above, we will need to develop some machinery. This will be essentially the first part of this course, that shall also include some other interesting applications.

In the second part we will discuss some topics related to game theory. One goal is to show a proof of Brouwer's fixed point using Sperner's Lemma, and then learn and understand how it relates to von Neumann's Minimax Theorem and the Nash Equilibrium of games. Then, we will return to the topics of the first part and show a lemma due to Scarf and its applications to combinatorics, also discussing some of its applications to game theory.

Finally, in the third part of this course, we shall prove Tucker's Lemma and see several of its applications, including the Borsuk-Ulam Theorem.

[1] It is easy but you should try nonetheless.

Among its combinatorial applications, arguably the most famous is the following.

A *colouring* of a graph is a partition of its vertex set into independent sets (you can think of each one as being coloured with a different colour). The minimum number of colours needed to colour a graph is called the *chromatic number* of G, denoted by $\chi(G)$. Given an arbitrary graph, the problem of computing $\chi(G)$ is generally hard, and has motivated most of the of the early development of graph theory.

The *Kneser graph* $K_{n:k}$ is the graph whose vertex set correspond to all k-subsets of a set with n elements, two of which are made adjacent if and only if they are disjoint. Note that $K_{n:k}$ is trivially a copy of the empty graph whenever $k > n/2$, and if $k = n/2$, then it is the disjoint union of edges.

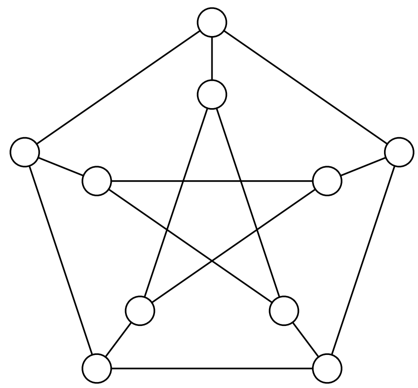

The Petersen Graph is a Kneser Graph. Who is n and k? Label the vertices correctly.

In 1955, Kneser conjectured that the chromatic number of $K_{n:k}$ is $n - 2k + 2$ provided $n \geq 2k$. It was only 1978 that Lovász proved this result true. The methods he used were based on Tucker's Lemma, and gave rise to the field of topological combinatorics.

2 Geometric simplicial complexes*

A set of vectors in $\mathbb{R}^d$ is *affinely dependent* if they are linearly dependent with the extra property that the coefficients of a linear combination giving zero also sum to zero. That is, $\{v_0, ..., v_k\}$ are affinely dependent if there are $\alpha_0, ..., \alpha_k \in \mathbb{R}$, not all zero, such that:

(1) $\alpha_0 v_0 + ... + \alpha_k v_k = 0$

(2) $\alpha_0 + ... + \alpha_k = 0$

A set of vectors is called *affinely independent* if they are not affinely dependent.

If $k = 1$, two affinely dependent vectors are equal. If $k = 2$, three three affinely dependent vectors are on the same line. If $k = 3$, then four affinely dependent vectors are on the same plane. In fact, no set of four vectors can be linearly independent in $\mathbb{R}^3$. However any four vectors not lying on a plane are affinely independent.

You should be able to show the following exercise.

Exercise 2.1. Vectors $v_1, ..., v_k$ in $\mathbb{R}^d$ are affinely dependent if and only if any of the following holds:

(1) $\{v_2 - v_1, ..., v_k - v_1\}$ are linearly dependent in $\mathbb{R}^d$.

(2) $\{(1, v_1), ..., (1, v_k)\}$ are linearly dependent in $\mathbb{R}^{d+1}$.

The *convex hull* of points $\{v_1, ..., v_k\}$ is the set of all points which are convex combinations of these, that is, the set

$$\{\alpha_1 v_1 + ... + \alpha_k v_k : \alpha_i \geq 0 \text{ and } \alpha_1 + ... + \alpha_k = 1\}.$$

A *simplex* σ is the convex hull of a set of $(k+1)$ affinely independent points $A = \{v_0, ..., v_k\}$ in $\mathbb{R}^d$. We say the points in A are the vertices of σ, and that the dimension of σ is k. A face of σ is the convex hull of any subset of points. Note that σ is a face of itself.

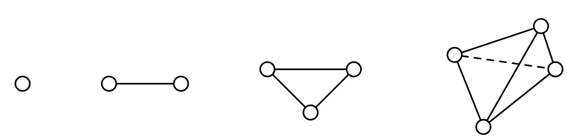

0-, 1-, 2-, and 3- simplices.

The *relative interior* of a simplex σ is the set of points that do not belong to any proper face, that is

$$\sigma - \bigcup\{\tau : \tau \text{ is a proper face of } \sigma\}.$$

Observe that any point in a simplex σ is in the relative interior of exactly one face of the simplex.

A set Σ of simplices is called a *geometric simplicial complex* if both conditions below hold.

- If $\sigma \in \Sigma$, then every face of σ is in Σ; and

- for all $\sigma, \tau \in \Sigma$, $\sigma \cap \tau$ is a face of both σ and τ (this allows $\sigma \cap \tau = \emptyset$.)

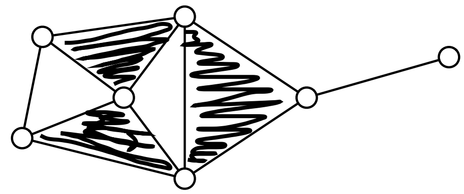

A simplicial complex. Simplices are the points, lines, and the shaded triangles.

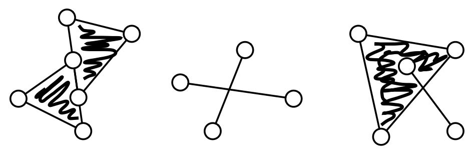

Not simplicial-complexes. Why?

The *dimension* of a simplicial complex is the maximum dimension of its simplices. Observe that a 1-dim complex in $\mathbb{R}^2$ is precisely a straight-line planar embedding of a graph (planar in the usual sense).

The *polyhedron* of Σ, denoted by $\|\Sigma\|$, is the union of its simplices: $\bigcup\{\sigma : \sigma \in \Sigma\}$. For all $k \leq \dim(\Sigma)$, the subcomplex of Σ consisting of all r-dim simplices where $r \leq k$ is the k-skeleton of Σ. For example, the 0-skeleton consists of all vertices in all $\sigma \in \Sigma$ (except for $\emptyset$). We call it the vertex set of Σ. We also write $\Sigma^{\leq k}$ for the k-skeleton of Σ.

Observe that for a simplicial complex Σ, the relative interior of all simplices in Σ partition $\|\Sigma\|$. For each $x \in \|\Sigma\|$, we call the unique σ that contains x in its relative interior the support of x, and denote it by $\text{supp}(x)$.

Finally, note that a simplex $\sigma \in \mathbb{R}^d$ along with all its faces forms a simplicial complex.

Exercise 2.2. Prove this fact, verifying that both conditions that define a simplicial complex are satisfied.

Let X be a set of points in $\mathbb{R}^d$. A geometric simplicial complex Σ in $\mathbb{R}^d$ is a *triangulation* of X if $\|\Sigma\|$ is homeomorphic to X (this basically means that you should be able to deform $\|\Sigma\|$ so that it becomes equal to X, but you cannot create or destroy holes). This allows us to talk about triangulation of balls or tori, even though the lines will not be straight.

The *boundary* of a d-simplex is the union of all faces which are $(d-1)$-simplices. This coincides with the set of points which are not in its relative interior. The boundary of σ is denoted by $\partial\sigma$.

Note that if Σ is a triangulation of a d-simplex in $\mathbb{R}^d$ and τ is a $(d-1)$-simplex of Σ that is not contained in $\partial\sigma$, then τ is a proper face of exactly two d-simplices in Σ. If τ is in $\partial\sigma$, then it is a proper face of exactly one d-simplex in Σ.

3 Abstract simplicial complexes*

An *abstract simplicial complex* is a set $\mathcal{A}$ of subsets of a finite set V such that if $A \in \mathcal{A}$ and $B \subseteq A$, then $B \in \mathcal{A}$. The sets $A \in \mathcal{A}$ are called simplices, and the dimension of $A \in \mathcal{A}$ is $|A| - 1$. For example, if Σ is a geometric simplicial complex, then take all sets of vertices of the simplices in Σ, and they form an abstract simplicial complex. Unless $\mathcal{A} = \emptyset$, then $\emptyset \in \mathcal{A}$, and is the unique simplex in $\mathcal{A}$ of dimension -1. We can see the ground set V as the set of 0-simplices of $\mathcal{A}$.

We say that $\mathcal{A}$ is *pure* if every maximal simplex in $\mathcal{A}$ has the same dimension d. We denote a pure simplicial complex of dimension d by d-PSC. The independent sets of matroids are a special class of pure simplicial complexes. The union of the vertex set with the edge set of a graph is a 1-PSC (as long as there are no isolated vertices).

Note that a d-PSC $\mathcal{A}$ is determined by the set $\mathcal{A}^d$ of all d-dimensional simplices in $\mathcal{A}$. Here $\mathcal{A}^d$ is a $(d+1)$-uniform hypergraph on V, meaning elements are $(d+1)$-subsets of V. We say $\mathcal{A}$ is the *(downward) closure* of $\mathcal{A}^d$.

The *boundary* of a d-PSC $\mathcal{A}$ is the closure of

$$\big\{B \subseteq V : \quad |B| = d, \quad |\{A \in \mathcal{A}^d : B \subseteq A\}| \equiv 1 \pmod 2\big\},$$

we write $\partial\mathcal{A}$ for the boundary. Note that the boundary of a graph (seen as a 1-PSC) is the set of vertices of odd degree.

If $\partial\mathcal{A} = \emptyset$, then we say that $\mathcal{A}$ is a d-sphere. For example, a graph where all vertices have even degree is a 1-sphere. A 0-sphere is a set of vertices of even size, as $\emptyset$ will be contained in an even number of maximal simplices.

The handshake lemma in Graph Theory states that the sum of the vertex degrees is equal to twice the number of edges — thus, the number of vertices of odd degree is even. We will see in the lemma below how this observation significantly generalizes for abstract simplicial complexes in terms of their boundary.

There is a convenient way of finding $\partial\mathcal{A}$. Define a matrix $\mathcal{A}[d, d+1]$ with rows indexed by d-subsets of V, and columns indexed by the $(d+1)$-subsets in $\mathcal{A}$, that is $\mathcal{A}^d$, having the (B, A) entry equal to 1 if $B \subseteq A$, 0 otherwise. By multiplying this matrix by the all 1s vector in $\mathbb{Z}_2$, we will obtain the characteristic vector of a d-uniform hypergraph. Its closure is the boundary of $\mathcal{A}$.

Lemma 3.1. *Let $\mathcal{A}$ be a d-PSC. Then $\partial\mathcal{A}$ is a $(d-1)$-sphere (the boundary of a boundary is empty.)*

Proof. Let $v = \mathcal{A}[d, d+1] \cdot \mathbf{1}$, then v is the characteristic vector of $(\partial\mathcal{A})^{d-1}$. Let $\mathcal{K}$ be the $(d-1)$-PSC that is the closure of the complete d-uniform hypergraph on $V(\mathcal{A})$, that is, $\mathcal{K}$ is the collection of all subsets of V with at most d elements. Then the boundary of $\partial\mathcal{A}$ is the closure of the $(d-1)$-uniform hypergraph whose characteristic vector is:

$$\mathcal{K}[d-1, d]v$$

But this is:

$$\mathcal{K}[d-1, d]\mathcal{A}[d, d+1]\mathbf{1}$$

The product of the matrices is counting the number of d-subsets that are sandwiched between a $d-1$ subset and a particular $d+1$ subset. This is either 0 or 2, so the matrix product is the 0 matrix in $\mathbb{Z}_2$ and thus this whole product is the 0-vector in $\mathbb{Z}_2$. Therefore $\partial\partial\mathcal{A}$ is empty and $\partial\mathcal{A}$ is a $(d-1)$-sphere. □

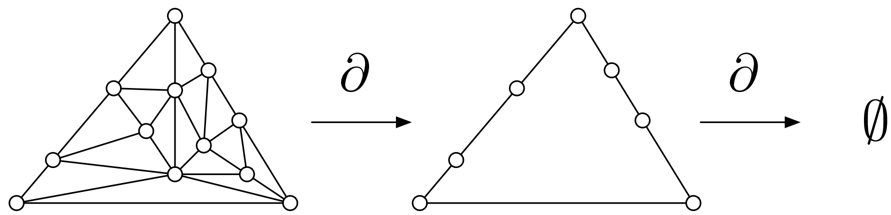

See the triangulation as a 2-PSC. The boundary of the boundary is empty.

Proposition 3.2. *Let $\mathcal{A}$ be a $(d-1)$-sphere, and w be a vertex not in $V(\mathcal{A})$. Then the closure $\mathcal{C}$ of $\{A \cup w : A \in \mathcal{A}^{d-1}\}$ is a d-PSC with boundary $\mathcal{A}$.*

Proof. Clearly $\mathcal{C}$ is a d-PSC.

Consider $A \in \mathcal{A}$ a $(d-1)$-simplex. Then A is contained in precisely one d-simplex of $\mathcal{C}$, namely, $A \cup w$. Therefore $\partial\mathcal{C} \supseteq \mathcal{A}$.

Suppose C is in $\partial\mathcal{C}$. If $w \in C$, then $C \backslash w$ is contained in an even number of $(d-1)$-simplices of $\mathcal{A}$ because $\mathcal{A}$ is a sphere, therefore C is contained in an even number of d-simplices of $\mathcal{C}$, hence it cannot be in the boundary of $\mathcal{C}$. So $w \notin C$, and hence there is a unique possible d-simplex containing C, which is $B = C \cup w$. By definition of the maximal simplices of $\mathcal{C}$, C must have been a simplex of $\mathcal{A}$, hence $\partial\mathcal{C} \subseteq \mathcal{A}$. □

4 Link and hole-filling*

Let $\mathcal{A}$ be a d-PSC and let $U \subseteq V(\mathcal{A})$ be a subset of size at most d. The *link* of U in $\mathcal{A}$ is:

$$\mathcal{L}_{\mathcal{A}}(U) = \{A \backslash U : A \in \mathcal{A}, U \subseteq A\}$$

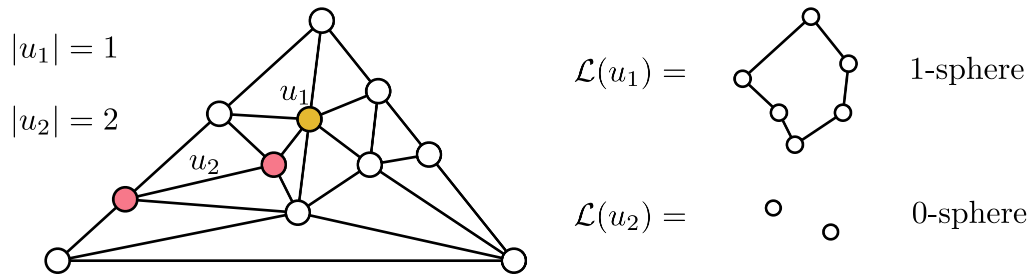

The link of two simplices in a triangulation.

In a graph (with no isolated vertices), the link of a vertex is its neighbourhood (plus $\emptyset$).

Lemma 4.1. *Let $\mathcal{A}$ be a d-PSC and let U be a simplex in $\mathcal{A}$ of size $u \leq d$ that is not a simplex of $\partial\mathcal{A}$. Then the link of U in $\mathcal{A}$ is a $(d-u)$-sphere.*

Proof. The maximal simplices of $\mathcal{L}(U)$ are precisely maximal simplices of $\mathcal{A}$ that fully contain U upon having U removed. These sets have size $d+1-u$. Consider a set B of size $d-u$ in V. If $B \cap U \neq \emptyset$, no simplex in $\mathcal{L}(U)$ contains B. Thus, assume otherwise. If B were contained in an odd number of maximal simplices of $\mathcal{L}$, then $B \cup U$ would be contained in an odd number of maximal simplices of $\mathcal{A}$. A contradiction to the hypotheses that U is not in the boundary of $\mathcal{A}$. So B is in an even number of maximal simplices of $\mathcal{L}(U)$, as wished. □

Lemma 4.2 (Hole filling)**.** *Let $\mathcal{B}$ be a $(d+1)$-PSC and let $B \in \mathcal{B}$ be a simplex that is not a simplex of $\partial\mathcal{B}$. Let $\mathcal{L} = \mathcal{L}_{\mathcal{B}}(B)$ and set $b = |B|$. Let $\mathcal{C}$ be a $(d-b+2)$-PSC with $\partial\mathcal{C} = \mathcal{L}$ and $V(\mathcal{C}) \cap B = \emptyset$.*

- *Suppose further that every $(d-b+2)$-simplex in $\mathcal{C}$ has a vertex that is not in $V(\mathcal{B})$.*

Let $\mathcal{B}'$ be the closure of

$$\mathcal{B}^{d+1} \setminus \{D \cup B : D \in \mathcal{L}, |D| = d-b+2\}$$
$$\cup \{C \cup B' : C \in \mathcal{C}, |C| = d-b+3, B' \subseteq B, |B'| = b-1\}$$

Then $\partial\mathcal{B}' = \partial\mathcal{B}$

Proof. Let A be a $(d+1)$-set. We want to show that the number of $(d+1)$-simplices in $\mathcal{B}'$ that contain A has the same parity as the number of $(d+1)$-simplices in $\mathcal{B}$ that contain A.

Let

$$S_1 = \{D \cup B : D \in \mathcal{L}, |D| = d-b+2, A \subseteq D \cup B\}$$

and

$$S_2 = \{C \cup B' : C \in \mathcal{C}, |C| = d-b+3, B' \subseteq B, |B'| = b-1, A \subseteq C \cup B'\}.$$

Note that $S_1 \subseteq \mathcal{B}^{d+1}$. Note that by definition of $\mathcal{C}$, $\mathcal{B} \cap S_2 = \emptyset$. So in order to prove that A is in the boundary of $\mathcal{B}'$ if and only if it is in the boundary of $\mathcal{B}$, it is enough to prove that $|S_2| - |S_1|$ is even. We split into cases:

(i) $B \subseteq A$. Note that every element of S_2 contains A, and so they should contain B as well. But the elements of S_2 are of the form $C \cup B'$, and in order for B to be in such set, $B \cap C \neq \emptyset$, contradicting one of the hypotheses. Thus $S_2 = \emptyset$. Now $|A \backslash B| = d - b + 1$, and $|S_1|$ counts exactly the number of $D \in \mathcal{L}$ containing such set. But $\mathcal{L}$ is a $(d - b + 1)$-sphere (by the previous lemma), so $|S_1|$ is even.

(ii) $|A \cap B| = b - 1$. Say $B_1 = A \cap B$. Note that $|A \backslash B| = d - b + 2$. If $D = A \backslash B$ and $D \in \mathcal{L} = \partial \mathcal{C}$, then the only element in S_1 is $D \cup B$, thus $|S_1| = 1$. But then there is an odd number of simplices in $\mathcal{C}$ containing D. If $D \notin \mathcal{L}$, then $|S_1| = 0$ but then the number of simplices in $\mathcal{C}$ containing D is even.

(iii) $|A \cap B| = b - 2$. Assume y and z are elements of B not in A. Note now that necessarily $|S_1| = 0$. Suppose $S_2 \neq \emptyset$, say $C \cup B' \in S_2$ contains A. As B' contains a vertex not in A, say y, this will be the only element of $C \cup B'$ not in A. So $C = A \backslash B$. But in this case there is exactly one other set B'' such that $A \subseteq C \cup B''$ with $C \cup B'' \in S_2$, namely $B'' = (A \cap B) \cup z$, and thus $|S_2|$ is even.

(iv) $|A \cap B| \leq b - 3$. Then both S_1 and S_2 are empty.

□

For example, let G be a graph with no isolated vertices. Consider u a vertex of even degree. The link of u is its neighbourhood $\Gamma(u)$. Take a graph C whose vertices of odd degree are exactly $\Gamma(u)$. Then the graph with edge set $E(G) - E_u \cup E(C)$ (plus the vertices on those edges) has the same set of odd degree vertices as G.

More interestingly, let $\mathcal{T}$ be a triangulation of a ball in $\mathbb{R}^3$. Let vw be a 1-simplex in $\mathcal{T}$ no contained in the boundary.

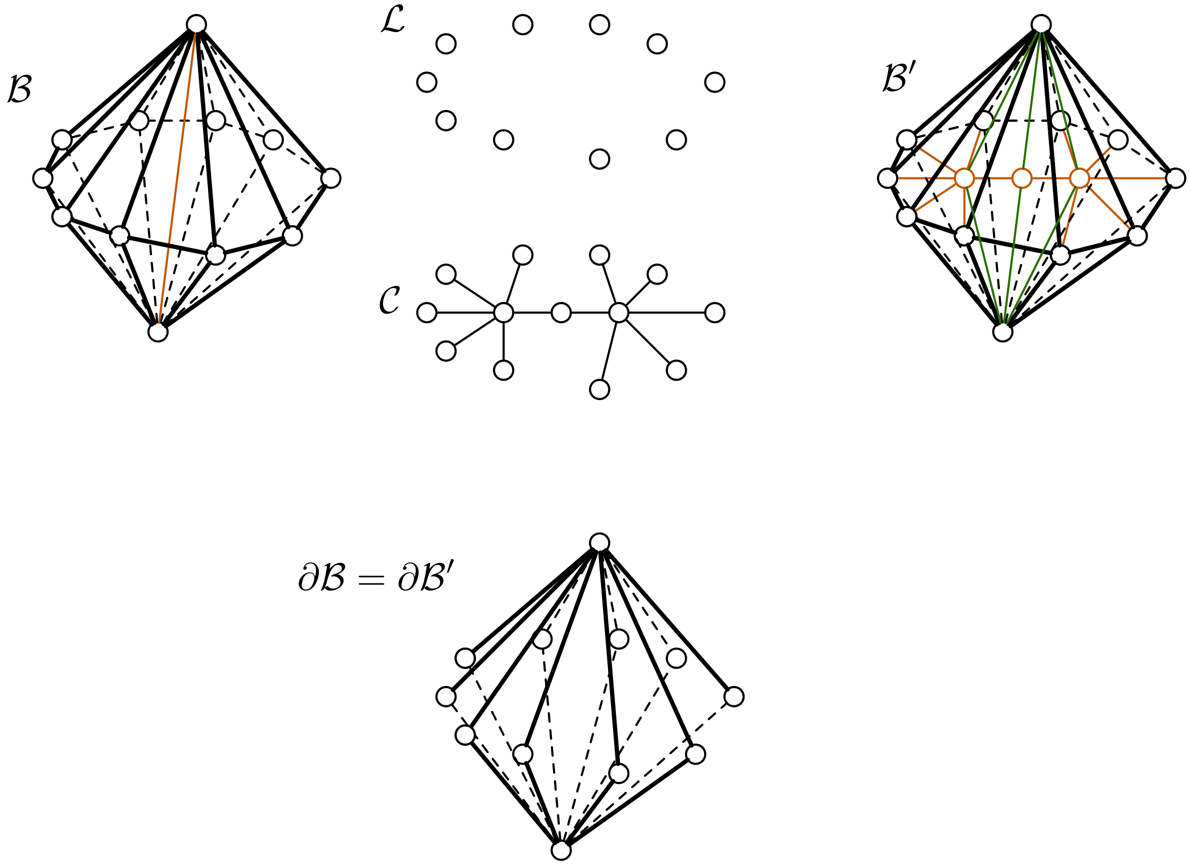

The hole-filling lemma.

5 Sperner's Lemma*

Lemma 5.1 (Sperner)**.** *Let* $S = \{v_1, ..., v_m\}$*. Suppose that for each* $T \subseteq [m]$*, there exists a* $(|T| - 1)$*-PSC* $\mathcal{A}_T$ *such that*

(i) $\mathcal{A}_\emptyset = \{\emptyset\}$ *and* $\mathcal{A}^0_{\{i\}} = \{\{v_i\}\}$*, and*

(ii) for each T*,* $|T| \geq 1$*, we have:*

$$\partial \mathcal{A}_T = \bigcup_{\substack{U \subseteq T \\ |U| = |T| - 1}} \mathcal{A}_U.$$

Suppose further that each vertex of $\mathcal{A}_{[m]}$ *is coloured with an element of* $[m]$ *such that for each* $T \subseteq [m]$*, if* $\{x\} \in \mathcal{A}_T$*, then the colour of* x *is in* T*. Then* $\mathcal{A}_{[m]}$ *contains a multicoloured* $(m-1)$*-simplex, that is, all colours* $\{1, ..., m\}$ *appear in the vertices of this simplex.*

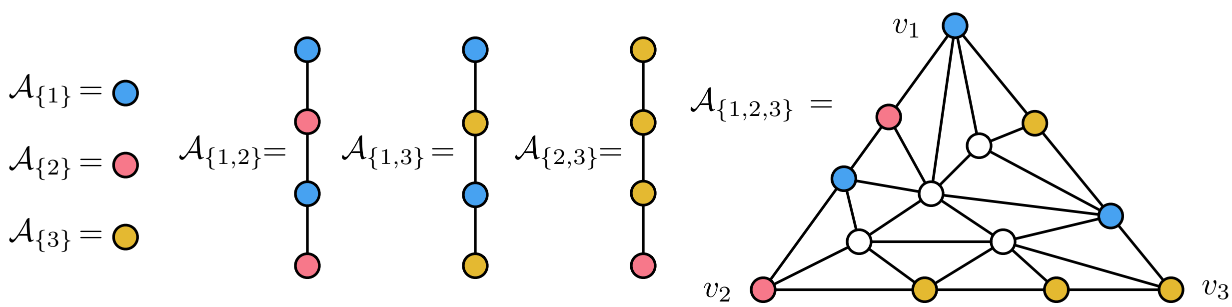

No matter how the remaining vertices of $\mathcal{A}_{\{1,2,3\}}$ are coloured, there will be a 3-chromatic simplex.

Proof. The proof is by induction on m. If $m = 1$, then $\{v_1\} \in \mathcal{A}_1$ is multicoloured. Assume $m \geq 2$ and that the number of multicoloured simplices for smaller m is odd.

Let $\mathcal{D}$ denote the set of $(m-2)$-simplices in $\mathcal{A}_{[m]}$ that are multicoloured with $[m-1]$. Note that any $(m-2)$-simplex of $\mathcal{D}$ in $\partial\mathcal{A}_{[m]}$ must be in $\mathcal{A}_{[m-1]}$. We apply induction, and so the number of $(m-2)$-simplices in $\mathcal{A}_{[m-1]}$ multicoloured with $[m-1]$ is odd.

Let H be a graph with vertex set $\mathcal{D}$, in which D and D' are joined by an edges if there exists an $(m-1)$-simplex $A \in \mathcal{A}_{[m]}$ that contains both. Observe that such A cannot be multicoloured by $[m]$.

Simplices of the boundary have degree 1, and there is an odd number of them. So there must be a simplex inside of degree 1, but this implies the existence of the required multicoloured simplex. More specifically:

Let $\text{mc}(D)$ be the number of multicoloured $(m-1)$-simplices in $\mathcal{A}_{[m]}$ that contain D. Let $\text{t}(D)$ be the total number of maximal simplices containig D. Then $\deg_H(D) = \text{t}(D) - \text{mc}(D)$. Then the number M of multicoloured simplices of $\mathcal{A}_{[m]}$ is:

$$M = \sum_{D \in \mathcal{D}} \text{mc}(D) = \sum_{D \in \mathcal{D}} (\text{t}(D) - \deg_H(D))$$

The sum of degrees is even, and by definition of boundary, we have

$$\sum_{D \in \mathcal{D}} \mathrm{t}(D) \equiv \sum_{D \in \mathcal{D} \cap \partial \mathcal{A}_{[m]}} \mathrm{t}(D) \pmod 2.$$

By induction hypothesis, it follows that M is odd. □

We should point out that the proof above is algorithmic, containing an implicit procedure to find the multicoloured simplex.

In order to apply Sperner's Lemma, the following technical lemma will be useful.

Lemma 5.2. *Let $S = \{v_1, ..., v_m\}$, $m \geq 2$, so that for each $T \subseteq [m]$ with $|T| \leq m-1$, there exists a $(|T|-1)$-PSC $\mathcal{A}_T$ such that*

(i) $\mathcal{A}_\emptyset = \{\emptyset\}$ *and* $\mathcal{A}^0_{\{i\}} = \{\{v_i\}\}$,

(ii) for each T, $|T| \geq 1$, we have:

$$\partial \mathcal{A}_T = \bigcup_{\substack{U \subseteq T \\ |U| = |T|-1}} \mathcal{A}_U,$$

(iii) for each $T \neq T'$ with $|T| = |T'|$, we have $\mathcal{A}_T^{|T|-1} \cap \mathcal{A}_{T'}^{|T|-1} = \emptyset$.

Then

$$\mathcal{A} = \bigcup_{\substack{T \subseteq [m] \\ |T| = m-1}} \mathcal{A}_T$$

is an $(m-2)$-sphere.

Proof. Consider $v = \mathcal{A}[m-2, m-1]\mathbf{1}$. Note that (iii) implies that $\mathcal{A}_T^{|T|-1}$, for $|T| = m-1$, is a partition of the $(m-1)$-sets which are simplices of $\mathcal{A}$. Hence, for an $(m-2)$-set B, the B-entry of v is just the sum of the B entries of $\mathcal{A}_T[m-2, m-1]\mathbf{1}$ over all T.

If B is not on the boundary of a $\mathcal{A}_{T_0}$, this entry is zero. Otherwise, B is on the boundary and there would be exactly one $U_0 \subseteq T_0$ such that $B \in \mathcal{A}_{U_0}$ (again, because of (iii)). But then there is another $T_1 \supseteq U_0$ such that $B \in \partial \mathcal{A}_{T_1}$. This proves that the B-entry of v is zero, and $\mathcal{A}$ is a sphere. □

The picture here is exactly that of a triangulation of a triangle.

6 Labelling and connectedness*

Let H be a graph. The *independence complex* $\mathcal{I}(H)$ of H has vertex set $V(H)$ and simplices the independent sets of H (note that subsets of independent sets are independent). In order to apply Sperner's lemma to a complex such as this, we will attempt to construct a nice pure simplicial complex $\mathcal{A}$ related to $\mathcal{I}$, and a map from $\mathcal{A}$ to $\mathcal{I}$ so that applying Sperner's lemma to $\mathcal{A}$ simulates applying it to $\mathcal{I}$.

Let Σ be a simplicial complex, and let $\mathcal{A}$ be a d-PSC. A *labelling* of $\mathcal{A}$ from Σ (also said a Σ-labelling) is a map $f : V(\mathcal{A}) \to V(\Sigma)$. A simplex $A \in \mathcal{A}$ is *properly labelled* by f if $\{f(a) : a \in A\} \in \Sigma$.

We call f a *proper labelling* of $\mathcal{A}$ if each $A \in \mathcal{A}$ is properly labelled by f.

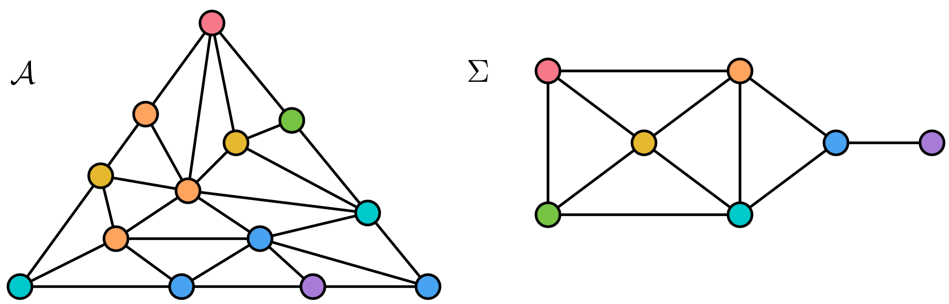

A proper Σ-labelling of the triangulation $\mathcal{A}$.

We say Σ is (topologically) k-connected if for each d, $-1 \leq d \leq k$, for every d-sphere $\mathcal{A}$ and for every proper Σ-labelling f of $\mathcal{A}$, there exists a $(d+1)$-PSC $\mathcal{B}$ and a proper Σ-labelling f' of $\mathcal{B}$ such that $\partial\mathcal{B} = \mathcal{A}$ and $f'|_{\mathcal{A}} = f$.

Exercise 6.1. Try to decide whether or not Σ in the example above is topologically 0-connected. How about 1-connected?

Note that a (-1)-connected simplicial complex is exactly those which are non-empty. Let G be a graph seen as a simplicial complex of dimension 1. Suppose G is 0-connected. Two vertices of G form a 0-sphere. A PSC $\mathcal{B}$ as in the definition is just a path between the two vertices, and finally any labelling will work. So this translates to the graph begin connected in the usual sense.

Lemma 6.2. *Let $\mathcal{A}$ be a d-sphere and let f be a proper Σ-labelling of $\mathcal{A}$. Suppose $\mathcal{B}$ is a $(d+1)$-PSC such that there exists a proper Σ-labelling f' of $\mathcal{B}$ with $\partial\mathcal{B} = \mathcal{A}$ and $f'|_{\mathcal{A}} = f$. Then there exists a $(d+1)$-PSC $\mathcal{B}_1$ and a proper Σ-labelling f_1 such that $\partial\mathcal{B}_1 = \mathcal{A}$ and $f_1|_{\mathcal{A}} = f$, and every $(d+1)$-simplex of $\mathcal{B}_1$ has a vertex that is not in $\partial\mathcal{B}_1$.*

Proof. Suppose $B \in \mathcal{B}^{(d+1)}$ has all its vertices in ∂B. Note that B defines a $(d+1)$-PSC $\mathcal{C}$ by $\mathcal{C}^{d+1} = \{B\}$. By Lemma 3.1, $\partial\mathcal{C}$ is a d-sphere. Let w be a vertex not in $V(B)$. By Proposition 3.2, the closure of $\{B \cup w : B \in \partial\mathcal{C}^d\}$ is a $(d+1)$-PSC with boundary $\partial\mathcal{C}$. Let $\mathcal{B}'$ be the closure of:

$$\mathcal{B}^{d+1} \backslash \{B\} \cup \{B' \cup w : B' \subseteq B, |B'| = |B| - 1\}.$$

Let A be a $(d+1)$-set. If $w \notin A$, then the number of $(d+1)$-simplices containing A in $\mathcal{B}'$ is the same as that of $\mathcal{B}$, because $A \subseteq B$ implies that A is one of the B'.

If $w \in A$, then $A \nsubseteq B$. If $A \subseteq B' \cup w$ for some B', then $A \subseteq B'' \cup w$ for some other B'' (because of the sizes). So the boundary of $\mathcal{B}'$ is the same as the one of $\mathcal{B}$.

The labelling is extended to $\mathcal{B}'$ by saying that $f'(w) = f(v)$ for some $v \in B$. Hence for every simplex $D \cup w \in \mathcal{B}$, we will have:

$$\{f'(v) : v \in D \cup w\} \subseteq \{f'(v) : v \in B\} \in \Sigma.$$

This process destroyed a bad simplex without creating any new one, so it can be applied iteratively for the desired result. □

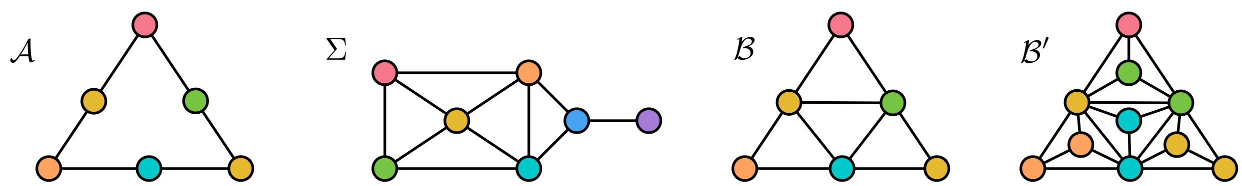

An example of Lemma 6.2.

Theorem 6.3. *Let Σ be a simplicial complex whose vertices are coloured from $[m]$. Suppose for each $T \subseteq [m]$, the subcomplex Σ_T (consisting of all simplices in Σ whose vertices are coloured from T) is $(|T|-2)$-connected. Then Σ contains a multicoloured $(m-1)$-simplex.*

Proof. The idea to prove this theorem is to apply Sperner's lemma. First note that the result is true for $m = 0$ and $m = 1$. Assume $m \geq 2$. To apply Sperner's lemma, we will construct, for each $T \subseteq [m]$, a $(|T|-1)$-PSC $\mathcal{A}_T$ such that

- There is a proper Σ-labelling f of $\mathcal{A}_{[m]}$.
- This labelling induces a colouring of c_f of $V(\mathcal{A}_{[m]})$ from $[m]$. More specifically, $c_f(v)$ is the colour in Σ of $f(v)$.
- $\{\mathcal{A}_T : T \subseteq [m]\}$ and the colouring c_f satisfy the conditions of Sperner's Lemma.

If we are able to do that, then, by Sperner's Lemma, $\mathcal{A}_{[m]}$ and hence Σ contains a multicoloured $(m-1)$-simplex.

To accomplish this task, we define:

- $\mathcal{A}_\emptyset = \{\emptyset\}$.
- Fix $\{v_1, ..., v_m\}$. Then set $\mathcal{A}_i^0 = \{\{v_i\}\}$.

We label v_i with some vertex of $\Sigma_{\{i\}}$ (which is non-empty by hypothesis — simply set $|T| = 1$ and use definition for -1-connectedness).

Now we fix T_0, $|T_0| \geq 2$, and suppose we have defined $\mathcal{A}_T$ for all $T \subseteq T_0$, $|T| < |T_0|$, and a labelling f of $\bigcup_{T \subseteq T_0} V(\mathcal{A}_T)$ such that:

(i) the restriction of f to each $V(\mathcal{A}_T)$ is a proper labelling of $\mathcal{A}_T$, and

(ii) $\{\mathcal{A}_T : T \subseteq T_0, |T| < |T_0|\}$ satisfies the conditions of the technical Lemma 5.2.

Observe that for $|T_0| = 2$, both hold. Also, by (ii) and Lemma 5.2,

$$\mathcal{C} = \bigcup_{\substack{T \subseteq T_0 \\ |T| = |T_0| - 1}} \mathcal{A}_T \quad \text{is a } (|T_0| - 2)\text{- sphere.}$$

and f is a proper labelling of $\mathcal{C}$ from $\bigcup_{T \subseteq T_0} \Sigma_T \subseteq \Sigma_{T_0}$.

By the hypothesis of the theorem, Σ_{T_0} is $(|T_0|-2)$-connected, so there is a $(|T_0|-1)$-PSC $\mathcal{A}_{T_0}$ with $\partial \mathcal{A}_{T_0} = \mathcal{C}$, and a proper Σ_{T_0}-labelling of $\mathcal{A}_{T_0}$ that extends f. Call the new labelling by f as well.

By Lemma 6.2, we may assume that every $(|T_0|-1)$-simplex in $\mathcal{A}_{T_0}$ has a vertex that is not in $\mathcal{C}$. We choose these to be completely new vertices, in the sense that each $(|T_0|-1)$-simplex in $\mathcal{A}_{T_0}$ has a vertex which is not in any $\mathcal{A}_{T'}$ with $|T'| = |T_0|$.

This completes the definition of $\mathcal{A}_{T_0}$ and f for each T_0. If T_1 has size $|T_0|+1$, then (i) and (ii) hold for it because of we constructed the new vertices in the T_0s. So the construction can be carried out to $\mathcal{A}_{[m]}$, where the hypotheses of Sperner's Lemma will be met. □

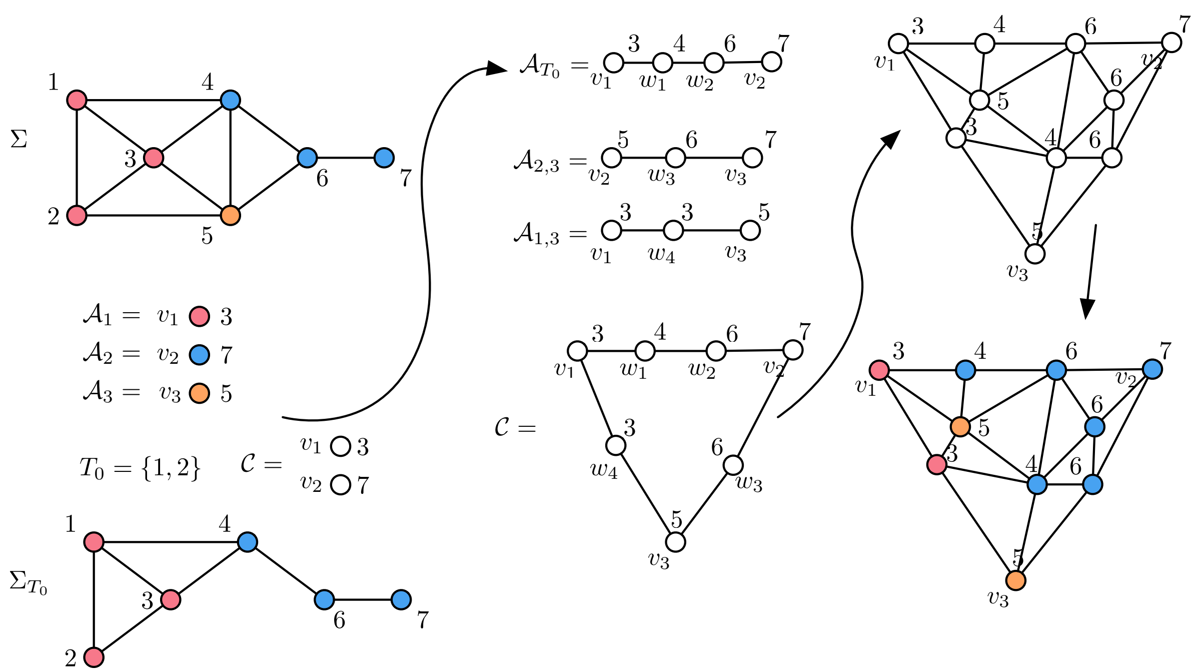

An example of Theorem 6.3.

7 Defect version of Theorem 6.3*

Our goal in this section is to walk towards a proof of a defect version of Theorem 6.3.

Theorem 7.1. *Let Σ be a simplicial complex whose vertices are coloured from $[m]$. Suppose Σ_T is $(|T|-d-2)$-connected for each $T \subset [m]$. Then Σ contains a multicoloured $(m-d-1)$-simplex.*

This will be eventually left as an exercise. We introduce below a construction and a useful lemma.

Let Σ be a simplicial complex whose vertices are coloured from $[m]$. Define a new simplicial complex as follows:

$$V(\Sigma^{(1)}) = V(\Sigma) \cup W^{(1)}$$

where $W^{(1)} = \{w_1, ..., w_m\}$ is disjoint from $V(\Sigma)$ and w_i has colour i. Then:

$$\Sigma^{(1)} = \{\sigma \cup \tau : \sigma \in \Sigma, \tau \in W^{(1)}, |\tau| \leq 1\}$$

We denote this by

$$\Sigma^{(1)} = \Sigma + W^{(1)}.$$

Lemma 7.2. *Let $d > 0$, let $T \subseteq [m]$, and suppose Σ_T is $(|T| - d - 2)$-connected. Then $\Sigma_T^{(1)}$ is $(|T| - d - 1)$-connected.*

Proof. If $|T| \leq d - 1$, there is nothing to prove. If $|T| = d$, then $\Sigma_T^{(1)}$ is non-empty, and therefore -1-connected. Assume $|T| \geq d + 1$. Let $1 \in T$.

Let r be so that $-1 \leq r \leq |T| - d - 1$, and let $\mathcal{A}$ be an r-sphere and f a proper labelling of $\mathcal{A}$ from $\Sigma_T^{(1)}$. Let x be a vertex not in $\mathcal{A}$ and consider $\mathcal{B}$ the closure of

$$\{A \cup x : A \in \mathcal{A}^r\}.$$

We know that $\mathcal{B}$ is an $(r + 1)$-PSC with boundary $\mathcal{A}$.

Set $f(x) = w_1$. If no vertex of $\mathcal{A}$ was labelled from $W^{(1)} \backslash w_1$, then f is a proper $\Sigma_T^{(1)}$-labelling of $\mathcal{B}$ and we are done.

Otherwise, there exists a non-empty set $\mathcal{B}_0$ of bad simplices in $\mathcal{B}$ whose label sets consist of a subset of $W^{(1)}$ of size equal 2 (because, after all, $\mathcal{A}$ was properly labelled). So each such label set is $\{w, w_1\}$ for some $w \in W^{(1)} \backslash w_1$, and each $B \in \mathcal{B}_0$ contains x. Let $B \in \mathcal{B}_0$ have largest dimension.

- If $B \in \mathcal{B}^{(r+1)}$, let y be a new vertex not in $V(\mathcal{B})$. Let $\mathcal{B}'$ be the closure of:

  $$\mathcal{B}^{(r+1)} \backslash B \cup \{B' \cup y : B' \subseteq B, |B'| = |B| - 1\}.$$

  We know that $\partial \mathcal{B}' = \partial \mathcal{B}$. Extend f by setting $f(y)$ to be any vertex in Σ_T different than w_1.

- If $\dim B \leq r$, we know, since B is not a simplex of $\mathcal{A} = \partial\mathcal{B}$, that the link $\mathcal{L}_{\mathcal{B}}(B)$ is an $(r-b+1)$-sphere, where $b = |B|$. Note that $\mathcal{L}$ is labelled by f from Σ_T (ie, no w_i appears as a label on $\mathcal{L}$) since B was a maximal bad simplex. Since $b \geq 2$ and Σ_T is $(|T|-d-2) \geq (r-b+1)$ -connected, there exists a $(r-b+2)$-PSC $\mathcal{C}$ whose boundary is $\mathcal{L}$ and an extension of the labelling $f|_{\mathcal{L}}$ to $\mathcal{C}$ that is a proper Σ_T-labelling. By the Lemma 5.2, we may assume that every $(r-b+2)$-simplex in $\mathcal{C}$ has a new vertex not in $V(\mathcal{B})$, so $V(\mathcal{C}) \cap V(\mathcal{B}) = \mathcal{L}$. Let $\mathcal{B}'$ be the closure of:

  $$\mathcal{B}^{(r+1)} \backslash \{D \cup B : D \in \mathcal{L}, |D| = r-b+2\} \cup \{C \cup B' : C \in \mathcal{C}^{(r-b+2)}, B' \subseteq B, |B'| = |B|-1\}.$$

  By the Hole Filling Lemma 4.2, $\partial\mathcal{B}' = \partial\mathcal{B}$.

In both cases, we obtain $\mathcal{B}'$ and an extended labelling f of $\mathcal{B}'$ from $\Sigma_T^{(1)}$ such that $\partial\mathcal{B}' = \partial\mathcal{B}$, $B \in \mathcal{B}'$ and no new bad simplices were created since added vertices were labelled from Σ_T. So we may repeat this construction to each $\mathcal{B}^*$. Hence $\Sigma_T^{(1)}$ is $(|T|-d-1)$-connected. □

Exercise 7.3. Prove Theorem 7.1.

Check the examples below on how to go from $\mathcal{B}$ to $\mathcal{B}^*$.

$d = 1$

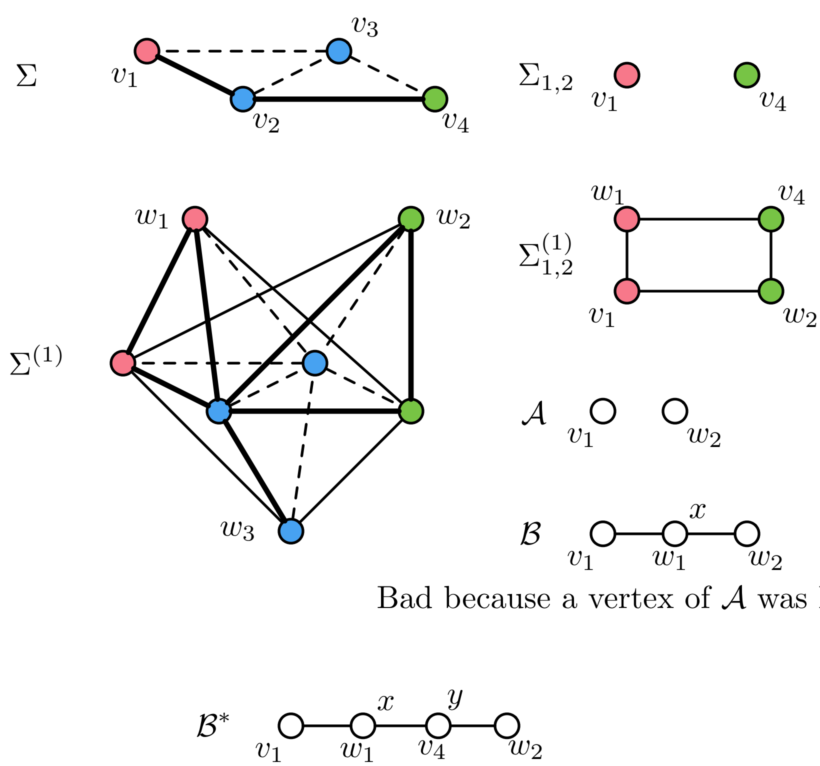

First case in the proof.

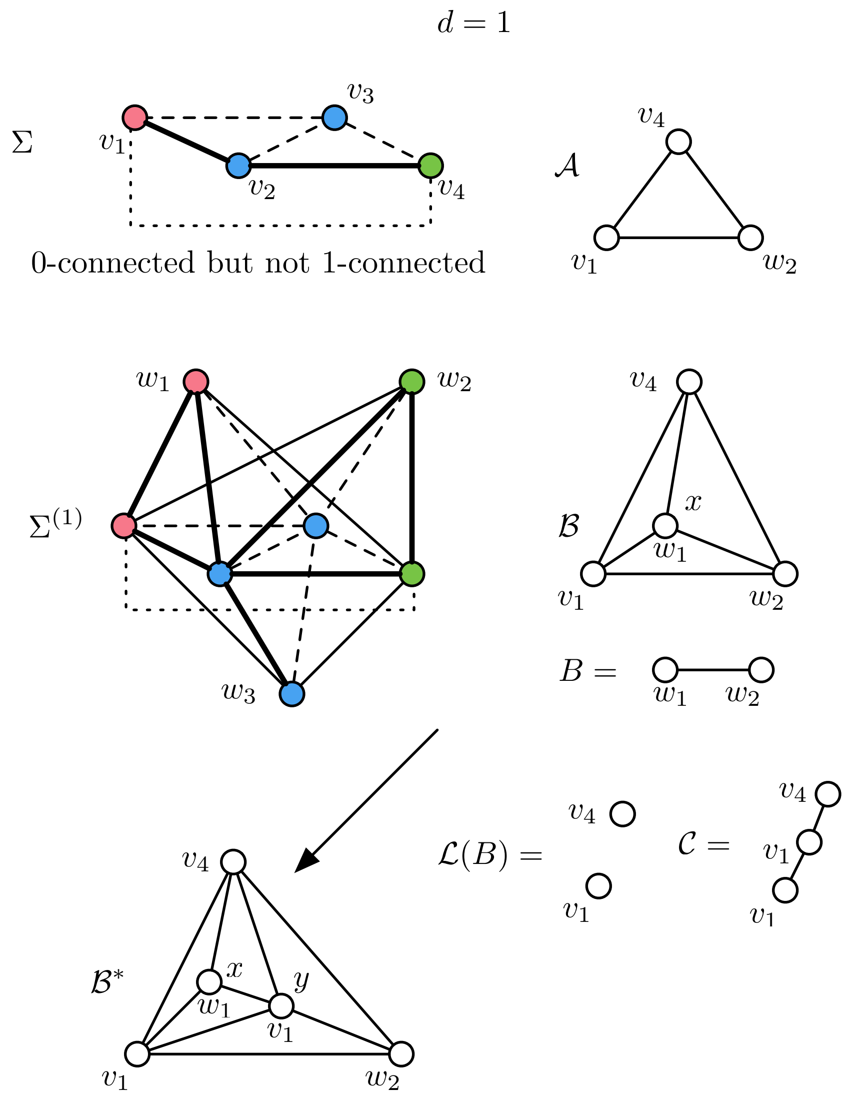

Second case in the proof.

8 Connectedness of $\mathcal{I}(G)$ via graph parameters*

Recall that one of our original motivating problems was that of finding an independent transversal in a graph. You can see a vertex partitioned graph is a graph whose $V(G)$ has been coloured — an independent transversal is therefore a multicoloured simplex of $\mathcal{I}(G)$. In order for us to be able to apply the technology developed in the previous sections, we need to be able to say something about the topological connectedness of $\mathcal{I}(G)$. This is the topic of this section.

Exercise 8.1. Show that if G has an isolated vertex, then $\mathcal{I}(G)$ is k-connected for every k.

Theorem 8.2. *Let $t \geq 0$, let G be a graph and let $xy \in E(G)$. Assume $\mathcal{I}(G \backslash xy)$ is t-connected, and $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$ is $(t-1)$-connected. Then $\mathcal{I}(G)$ is t-connected.*

Proof. Let $r \leq t$, and let $\mathcal{A}$ be an r-sphere properly $\mathcal{I}(G)$-labelled with f. Since $\mathcal{I}(G - xy)$ is t-connected, there exists an $(r+1)$-PSC $\mathcal{B}$ and a proper labelling f' of $\mathcal{B}$ from $\mathcal{I}(G - xy)$ with $\partial \mathcal{B} = \mathcal{A}$ and $f'|_{\mathcal{A}} = f$. By Lemma 6.2, we may assume each $(r+1)$-simplex of $\mathcal{B}$ has a vertex not in $\mathcal{A}$.

If no simplex of $\mathcal{B}$ has label set $\{x, y\}$, then f' is a proper $\mathcal{I}(G)$-labelling of $\mathcal{B}$, and we are done. Otherwise, let $\mathcal{B}_0$ be a the set of such bad simplices of $\mathcal{B}$.

Let $B \in \mathcal{B}_0$ be of largest dimension. Note that B is not a simplex of $\partial \mathcal{B} = \mathcal{A}$, since f was proper on $\mathcal{A}$. Again, we have two cases:

- If $B \in \mathcal{B}^{(r+1)}$, let $\mathcal{B}'$ be the closure of

  $$\mathcal{B}^{(r+1)} \backslash \{B\} \cup \{B' \cup w : B' \subseteq B, |B'| = |B| - 1\}.$$

  Then, as before, $\partial \mathcal{B}' = \partial \mathcal{B}$. Set $f'(w)$ to be a vertex of $G - (\Gamma(x) \cup \Gamma(y))$, which exists since $t \geq 0$ and $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$ is $(t-1)$-connected, hence non-empty.

- If $\dim B \leq r$, we know that $\mathcal{L} = \mathcal{L}_{\mathcal{B}}(B)$ is a $(r - b + 1)$-sphere ($b = |B|$). Then $\mathcal{L}$ is properly labelled from $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$ by f', since f' is a proper $\mathcal{I}(G \backslash xy)$-labelling of $\mathcal{B}$, and B is maximal. Since $b \geq 2$ and $r - b + 1 \leq t - 1$, and $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$ is $(t-1)$-connected, there exists an $(r - b + 2)$-PSC $\mathcal{C}$ with $\partial \mathcal{C} = \mathcal{L}$ and a proper $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$-labelling of $\mathcal{C}$ that extends f'. As usual, we may assume by Lemma 6.2 that each $(r - b + 2)$-simplex in $\mathcal{C}$ has a vertex not in $\partial \mathcal{C}$. So $V(\mathcal{C}) \cap V(\mathcal{B}) = \mathcal{L}$. Let $\mathcal{B}'$ be the closure of

  $$\mathcal{B}^{(r+1)} \backslash \{D \cup B : D \in \mathcal{L}^{(r-b+1)}\} \cup \{C \cup B' : C \in \mathcal{C}^{(r-b+2)}, B' \subseteq B, |B'| = b - 1\}.$$

  By the Hole Filling Lemma 4.2, $\partial \mathcal{B}' = \partial \mathcal{B}$, and f' is a proper $\mathcal{I}(G \backslash xy)$-labelling of $\mathcal{B}'$. Moreover, $B \notin \mathcal{B}'$, and no new bad simplices were introduced, since all new vertices were labelled from $G - (\Gamma(x) \cup \Gamma(y))$.

Upon repeating, we are able to find $\mathcal{B}^*$ and a proper $\mathcal{I}(G)$-labelling of $\mathcal{B}^*$ that extends f, thus $\mathcal{I}(G)$ is t-connected. □

Let G be a graph. A set of vertices U *(strongly) dominates* G if $x \in \Gamma(U)$ for each $x \in V(G)$.

Corollary 8.3. *Let G be a graph. Suppose G is not dominated by the vertices of any set F of edges of G with $|F| \leq t+1$. Then $\mathcal{I}(G)$ is t-connected.*

Proof. This is true for $t = -1$. Assume $t \geq 0$. We use induction on $|E(G)|$. If $|E(G)| = 0$, $\mathcal{I}(G)$ is k-connected for all k. Assume $|E(G)| \geq 1$, let $xy \in E(G)$.

If $\mathcal{I}(G)$ is not t-connected, then by the previous theorem either $\mathcal{I}(G - xy)$ is not t-connected, or $\mathcal{I}(G - (\Gamma(x) \cup \Gamma(y)))$ is not $(t-1)$-connected. In the former case, by induction, $G - xy$ is dominated by the vertices of a set of edges F with $|F| \leq t-1$, then G is also dominated by the same set. In the latter case, then add xy and we have the same contradiction. □

Corollary 8.4. *Suppose G is not dominated by a set of $\leq 2t+2$ vertices. Then $\mathcal{I}(G)$ is t-connected.* □

Corollary 8.5. *Let G be a graph and let $V_1 \cup \ldots \cup V_m$ be a partition of $V(G)$. If $|V_i| \geq 2\Delta(G)$ for all i, then G has an independent transversal with respect to this partition.*

Proof. For each $T \subseteq [m]$, consider $G_T = G[\bigcup V_i : i \in T]$. Number of vertices in G_T is at least $|T| \cdot 2\Delta$, so it cannot be dominated by $2(|T|-2)+2$ vertices, since each vertex has degree $\leq \Delta$. Hence $\mathcal{I}(G)$ is $(|T|-2)$-connected. By Theorem 7.1, $\mathcal{I}(G)$ has a multicoloured $(m-1)$-simplex, that is, and independent transversal. □

Let G be a graph. Let $i(G)$ be the largest ℓ such that there exists an independent set in G that is not dominated by any set of size $\leq \ell$ vertices.

For example, if S_n denotes de star graph with n leaves, and if P_n denotes the path graph with n vertices, it is immediate to verify that $i(S_n) = 0$, as every independent set is dominated by one vertex, and $i(P_n) = \lfloor n/2 \rfloor - 1$.

For a more interesting example, consider the cycle graph C_n with n odd. Its (strong) domination number is $\lfloor n/2 \rfloor$, but all of its independent sets can be dominated with $\lfloor n/2 \rfloor - 1$ vertices, hence $i(C_n) = \lfloor n/2 \rfloor - 2$.

Theorem 8.6. *Let G be a graph. If $i(G) \geq t+1$, then $\mathcal{I}(G)$ is t-connected.*

Proof. If $t = -1$, the result is trivially true. Assume $t \geq 0$. We apply induction on $|E(G)|$. If $|E| = 0$, we are done by Exercise 8.1. Assume $|E| \geq 1$. Let W be an independent set in G that is not dominated by $t+1$ vertices. Then $W \neq \emptyset$. Suppose $\mathcal{I}(G)$ is not t-connected. Let $w \in W$, then w is not an isolated vertex. Let v be a neighbour of w. We use Theorem 8.2 again. Note that, by induction, $\mathcal{I}(G - wv)$ is t-connected, thus it must be that $\mathcal{I}(G - (\Gamma(w) \cup \Gamma(v))))$ is not $(t-1)$-connected. By induction, this says that every independent set, in particular $W - \Gamma(v)$, is dominated by a set U of $\leq t$ vertices. We add v to U and we have a contradiction. □

9 Hall and Konig

In graph theory, there are two famous theorems about matchings in bipartite graph which play a central role in the theory and also were quite significant for the history of the field.

Let G be a graph. It is called bipartite if its vertex set can be partitioned into two classes and no edge contains two vertices of the same class. A matching of a graph is a collection of edges so that no two of them are incident to the same vertex.

Theorem 9.1 (Hall)**.** *Let G be a bipartite graph, with bipartition given by $A, B \subseteq V(G)$. Then there exists a a matching of G using all vertices in A if and only if, for all $S \subseteq A$, $|S| \leq |\Gamma(S)|$.*

It is easy to see how the condition is necessary to the existence of a matching for A. The somewhat surprising thing is that such clearly necessary condition is also sufficient.

A vertex cover of G is a subset $C \subseteq V(G)$ so that each edge of G is incident to at least one vertex of C. Let $\nu(G)$ denote the maximum size of a matching, and $\tau(G)$ denote the minimum size of a vertex cover. Clearly, if M is a matching of G and C is a vertex cover, at least one vertex of each edge in M must be in C. Thus, for all graphs, the inequality $\nu(G) \leq \tau(G)$ holds trivially. For bipartite graphs, the equality holds.

Theorem 9.2 (König)**.** *If G is a bipartite graph, then $\nu(G) = \tau(G)$.*

Both these theorems can be proved elementarily — if you have never seen these proofs, you are invited to research them on your own. Also, it is not difficult to prove Hall's Theorem using König's Theorem as a lemma. Finally, König's Theorem is a combinatorial manifestation of linear program strong duality (plus the fact that the incidence matrix of a graph is totally unimodular), making both results cornerstones in the field of combinatorial optimization.

In this course, however, we are interested in seeing how these results relate to topological methods — in fact we will be able to prove versions of these results to hypergraphs.

10 Combinatorial Applications*

A *hypergraph* $\mathcal{H}$ is a set of subsets (called edges) of a finite set $V(\mathcal{H})$. We say that $\mathcal{H}$ is r-uniform if each edge has size r. A bipartite hypergraph has vertex set $A \cup X$ such that each edge of $\mathcal{H}$ contains exactly one vertex of A. A *matching* is a set of disjoint edges.

Theorem 10.1 (Hall's Theorem for hypergraphs)**.** *Let $\mathcal{H}$ be a bipartite r-uniform hypergraph with vertex classes $A \cup X$. Suppose for each subset $T \subseteq A$, the neighbourhood $\Gamma_{\mathcal{H}}(T) = \{f \subseteq X : f \cup \{s\} \in \mathcal{H}$ for some $s \in T\}$ (which by the way is an $(r-1)$-uniform hypergraph) has a matching of size $\geq (r-1)(|T|-1)+1$. Then $\mathcal{H}$ has a matching of size $|A|$.*

Proof. Define the graph with vertex set $\{f_s \subseteq X : s \in A,\ f \cup s \in \mathcal{H}\}$. Note that $|f_s| = r-1$. Join f_s and $f'_{s'}$ if and only if $f_s \cap f'_{s'} = \emptyset$. Consider the partition of $V(G)$ where classes are $V_a = \{f_a : f_a \in V(G)\}$, for all $a \in A$. Observe that an A-perfect matching in $\mathcal{H}$ corresponds to an independent transversal of G with respect to the given partition. Let $T \subseteq [m]$, where $m = |A|$. Consider G_T. Then the condition implies that G_T has an independent set of size $\geq (r-1)(|T|-1)+1$. This is a set of disjoint $(r-1)$-subsets of X, which therefore cannot be dominated by $|T|-1$ vertices of G_T. So $i(G) \geq |T|-1$. By Theorem 8.6, $\mathcal{I}(G_T)$ is $(|T|-2)$-connected. Hence G has an independent transversal by Theorem 6.3. □

For a hypergraph $\mathcal{H}$, let $\nu(\mathcal{H})$ denote the maximum size of a matching in $\mathcal{H}$. A *cover* of $\mathcal{H}$ is a subset of vertices $C \subseteq V(\mathcal{H})$ such that every edge of $\mathcal{H}$ contains a vertex of C. We denote by $\tau(\mathcal{H})$ the minimum size of a cover of $\mathcal{H}$. We say an r-uniform hypergraph $\mathcal{H}$ is r-partite if $V(\mathcal{H}) = W_1 \cup W_r$ is a partition of V, and each edge of $\mathcal{H}$ has exactly one vertex in each W_i.

Theorem 10.2 (König's Theorem for hypergraphs)**.** *If $\mathcal{H}$ is a 3-partite, 3-uniform hypergrah, then*

$$\tau(\mathcal{H}) \leq 2\nu(\mathcal{H}).$$

Note that $\tau(\mathcal{H}) \leq r\nu(\mathcal{H})$ is easy to show for r-uniform hypergraphs. The result above is estabilishing $\tau(\mathcal{H}) \leq (r-1)\nu(\mathcal{H})$ for $r = 3$ (note that König's Theorem is this inequality with $r = 2$). For r is in general, this is inequality is a conjecture due to Ryser.

Lemma 10.3. *If $\mathcal{H}$ is an r-uniform bipartite hypergraph with vertex classes $A \cup X$, then for each $T \subseteq A$,*

$$\tau(\Gamma_{\mathcal{H}}(T)) \geq |T| - |A| + \tau(\mathcal{H}).$$

Proof. Note that for any cover C of $\Gamma_{\mathcal{H}}(T)$, $(A \backslash T) \cup C$ is a cover of $\mathcal{H}$. So, if $|C| = \tau(\Gamma_{\mathcal{H}}(T))$, we see that

$$\tau(\mathcal{H}) \leq |(A \backslash T) \cup C| = |A| - |T| + \tau(\Gamma_{\mathcal{H}}(T)).$$

□

Proof of theorem. Note that $\mathcal{H}$ is bipartite with vertex classes W_1 and $W_2 \cup W_3$, and for each $T \subseteq W_1$, $\Gamma_{\mathcal{H}}(T)$ is a bipartite graph. Hence, by the lemma with $A = W_1$, and by König's Theorem, we have, for each $T \subseteq A$,

$$\nu(\Gamma_{\mathcal{H}}(T)) \geq |T| - |A| + \tau(\mathcal{H}).$$

Define the graph G with $V(G) = \{f_s \subseteq X, f \cup s \in \mathcal{H}\}$, and $f_s f'_{s'} \in E(G)$ if and only if $f_s \cap f'_{s'} \neq \emptyset$, and partition its vertex set as $V_a = \{f_a : f_a \in V(G)\}$ for all $a \in A$. We want to find an independent transversal of G of size

$$\geq \left\lceil \frac{\tau(\mathcal{H})}{2} \right\rceil = |A| - \underbrace{\left(|A| - \left\lceil \frac{\tau(\mathcal{H})}{2} \right\rceil\right)}_{d} = |A| - d.$$

We therefore need to verify, by Theorem 7.1, that $\mathcal{I}(G_T)$ is $(|T| - d - 2)$-connected for each $T \subseteq A$.

Assume there exists $S_0 \subseteq A$ with $\mathcal{I}(G_{S_0})$ not $(|S_0| - d - 2)$-connected. Then, by Theorem 8.6, $i(S_0) \leq |S_0| - d - 2$. So every independent set in G_{S_0} is dominated by $\leq |S_0| - d - 1$ vertices in G_{S_0}. Thus

$$\nu(\Gamma_{\mathcal{H}}(S_0)) \leq 2(|S_0| - d - 1),$$

therefore $|S_0| - |A| + \tau(\mathcal{H}) \leq 2|S_0| - 2d - 2 \leq 2|S_0| - 2|A| + \tau(\mathcal{H}) + 1 - 2$. Thus $|A| + 1 \leq |S_0|$, but $S_0 \subseteq A$. This contradiction and Theorem 7.1 shows that

$$\nu(\mathcal{H}) \geq \left\lceil \frac{\tau(\mathcal{H})}{2} \right\rceil.$$

□

A third application is the related to a problem popularized by ErdǒS in the 80s. Let C_n be the cycle graph, and assume $3|n$. Add $n/3$ disjoint triangles on the vertex set $V(C_n)$. Is there an independent set in the resulting graph of size $n/3$?

Theorem 10.4. *Let G be a graph that is a disjoint union of cycles whose length are all divisible by 3. Let V_1,...,V_m be a partition of $V(G)$ where $|V_i| = 3$ for all i. Then G has an independent transversal.*

Proof. Let $S \subseteq [m]$. We will show that $i(G_S) \geq |S| - 1$ for all S. The graph G_S is a disjoint union of paths, and cycles whose length is divisible by 3. Let I_S be an independent set in G_S, consisting of every third vertex in each component. Then, since the cycle lengths are divisible by 3, $|I_S| \geq |V(G_S)|/3 \geq |S|$. Note that I_S cannot be dominated by $|S| - 1$ vertices. Hence $i(G_S) \geq |S| - 1$, and thus, by Theorem 8.6 and 6.3, G has an independent transversal. □

11 Fixed-point theorems

In this section, we start the second part of the course. Our goal now is to explore some applications of topological methods to game theory.

The most famous use of topology in game theory is arguably the proof that finite games admit mixed (probabilistic) strategies which attain a certain well defined notion of equilibrium — known as the Nash Equilibrium.

In order to show this result, we will now start from Sperner's Lemma and visit some famous results, most notably, Brouwer's fixed point theorem.

Because this course does not require students to have taken a first course in topology or analysis, I will go easy on some details.

Recall that we started our course talking about geometric simplicial complexes, and later moved to their abstract version, with which we proved Sperner's Lemma. Let us now start spelling Sperner's Lemma for geometric simplicial complexes.

For what follows, let Δ^r denote the r-dimensional regular simplex with edge length 1 and vertices $v_0, ..., v_r$, and let $F_0, ..., F_r$ denote its faces of dimension $r-1$ (also called *facets*). We assume

$$F_i = \text{conv}(\{v_j\}_{j \neq i}).$$

A *triangulation* $\mathcal{T}$ of Δ^r is a (pure) geometric simplicial complex of dimension r and whose union of its simplices is Δ^r. In other words, it is the downward closure a family of r-dimensional simplices whose union is Δ^r and so that the intersection of any two simplices in the family is a face of both.

For the lemma below, use the notation introduced above.

Lemma 11.1. *Assume vertices of $\mathcal{T}$ are coloured with colours $\{0, ..., r\}$. The following are equivalent.*

(i) No vertex of $\mathcal{T}$ which belongs to facet F_i received colour i.

(ii) For any vertex v contained in the face conv$(\{v_i\}_{i \in S})$ for $S \subseteq \{0, ..., r\}$, v is coloured with a colour from S.

Proof. Clearly (ii) implies (i). To see the converse, assume v is a vertex in $\text{conv}(\{v_i\}_{i \in S})$ for some S and v is badly coloured. Say with colour $j \notin S$. But v belongs to F_j, as

$$\text{conv}(\{v_i\}_{i \in S}) \subseteq \text{conv}(\{v_i\}_{i \neq j}) = F_j.$$

□

Note that Condition (ii) of the Lemma above (along with the definition of an abstract simplicial complex whose simplies are the vertex sets of the simplices of $\mathcal{T}$) is precisely the construction and the colouring that appears in the hypothesis of Sperner's Lemma. Therefore we have the result.

Lemma 11.2 (Sperner's lemma for geometric simplicial complexes). *Let $\mathcal{T}$ be a triangulation of Δ^r. Colour the vertices of $\mathcal{T}$ with colours from $\{0, 1, ..., r\}$ so that vertices on face F_i are never coloured with color i. Then there exists a multi-coloured simplex of dimension r in $\mathcal{T}$.* □

We will now meet the set covering version of Sperner's lemma, proved by Knaster, Kuratowski and Mazurkiewicz in 1929.

Lemma 11.3 (KKM 1929)**.** *Let $A_0, ..., A_r$ be closed sets contained in Δ^r. Assume $A_i \cap F_i = \emptyset$, and also that $A_0 \cup A_1 ... \cup A_r = \Delta^r$. Then*

$$A_0 \cap A_1 \cap ... \cap A_r \neq \emptyset.$$

Proof. For any $x \in \Delta^r$, let $d(x, A_i)$ denote the distance of x from A_i (this is well defined because each A_i is a closed set). Define the function

$$f(x) = \max\{d(x, A_i)\}_{i=0,...,r}.$$

Suppose the intersection of the A_is is empty. Then $f(x) > 0$ for all $x \in \Delta^r$, and because f is a continuous function, a theorem due to Weierstrass guarantees that the minimum value of f is attained within Δ^r, thus there exists $\varepsilon > 0$ so that $f(x) > \varepsilon$ for all $x \in \Delta^r$.

Now consider a triangulation of Δ^r so that the diameter of the r-simplices in $\mathcal{T}$ are at most ε, and colour each vertex v of this triangulation with an index j as long as $v \in A_j$. Note that Sperner's lemma hypothesis is met for this colouring, therefore there is a multi-coloured r-simplex S. But, then, for any $x \in S$, it follows that $f(x) \leq \varepsilon$. □

Much like the two equivalent conditions that can be used to state Sperner's Lemma, as seen in Lemma 11.1, there is also another equivalent condition that can be used instead in the KKM statement.

Proposition 11.4. *Let Δ^r be the r-dimensional simplex with vertices $v_0, ..., v_r$, and let $A_0, ..., A_r$ be closed sets contained in Δ^r and so that $A_0 \cup ... \cup A_r = \Delta^r$. The following two conditions are equivalent.*

(i) $A_i \cap F_i = \emptyset$ for all i.

(ii) For any $S \subseteq \{0, ..., r\}$,

$$conv(\{v_i\}_{i \in S}) \subseteq \bigcup_{i \in S} A_i.$$

Proof. Exercise. □

We are now ready to provide a proof of Brouwer's fixed point theorem for regular simplices. The full proof for compact convex sets requires the use of the notion of homeomorphisms (continuous bijections with continuous inverses). All compact convex subsets C of $\mathbb{R}^n$ are homeomorphic (to Δ^r). Thus any continuous functions from C to itself can be translated to a continuous function from Δ^r to itself, and the same goes for its fixed points.

Theorem 11.5 (Brouwer 1912)**.** *Let f be a continuous function from Δ^r to itself. Then there is $z \in \Delta^r$ so that $f(z) = z$, that is, f contains a fixed point.*

Proof. Suppose otherwise that f has no fixed points. Thus the function $g(x) = |x - f(x)|$ is continuous and is nowhere 0, therefore its minimum is attained in Δ^r and is positive. Let $\varepsilon > 0$ be so that $g(x) > \varepsilon$ for all $x \in \Delta^r$.

Let e_i denote the outward oriented normal unit vector of the facet F_i. Define the sets

$$A_i = \{x \in \Delta^r : \langle (f(x) - x), e_i \rangle \geq \varepsilon / r\}.$$

Note that

(a) Each A_i is closed.

(b) Their union is Δ^r. In fact, let $x \in \Delta^r$. The inequalities

$$\langle y, e_i \rangle \leq |f(x) - x|/r$$

define a regular simplex centred at the origin. Its inscribed ball has radius $|f(x) - x|/r$, thus this simplex is contained in a ball with radius $|f(x) - x|$. Thus, the point $y = f(x) - x$ is either outside this simplex or on its boundary. Therefore there is an i for which $\langle (f(x) - x), e_i \rangle \geq |f(x) - x|/r$, that is, $x \in A_i$ for at least some i.

(c) $A_i \cap F_i = \emptyset$, as $|f(x) - x| > 0$ for all x, but if $x \in F_i$, then $\langle (f(x) - x), e_i \rangle \leq 0$.

(d) $A_0 \cap ... \cap A_r = \emptyset$, because $\sum e_i = 0$, and thus $\sum \langle f(x) - x, e_i \rangle = 0$, so for at least one index, we have $\langle f(x) - x, e_i \rangle \leq 0$.

These contradict the KKM Theorem. □

Brouwer's fixed point theorem can be extended to a stronger result.

Theorem 11.6 (Kakutani's fixed point theorem)**.** *Let S be a compact convex set in $\mathbb{R}^n$, and let $f : S \to 2^S$ be a function that maps elements of S to subsets of S, such that*

(i) $f(x)$ is non-empty and convex for all $x \in S$.

(ii) if $(x_n, y_n)_{n \geq 0}$ converges to (x, y) with $y_n \in f(x_n)$, then $y \in f(x)$.

Then there exists $x \in S$ such that $x \in f(x)$.

Note that Brouwer's fixed point theorem is an immediate consequence of Kakutani's. Let us not prove this theorem. As before, we will restrict ourselves to the case where $S = \Delta^r$.

In order to follow this proof, we will need to introduce the concept of a *barycentric subdivision* of a simplex. This is obtained by adding a new vertex in the barycentre of each face of σ, and forming new faces which are the convex hulls of the barycentres. The barycentric subdivision of a triangulation is obtained upon replacing each simplex of the triangulation by its barycentric subdivision.

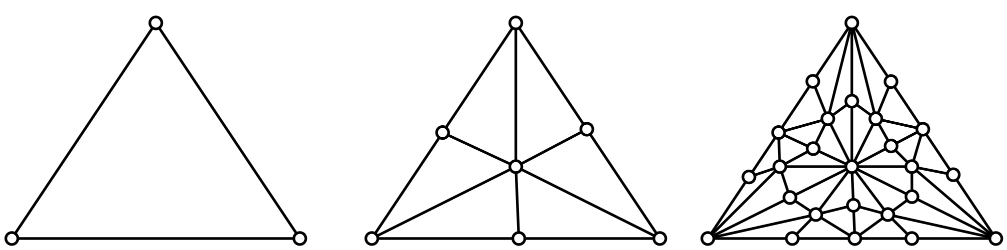

Note that the maximum diameter of a simplex in successive barycentric subdivision of a triangulation tends to 0 as long as you keep subdividing.

Proof. Let Σ_n be the n-th barycentric subdivision of Δ^r. Construct the function f_n as follows:

(a) For all v which is a vertex of Σ_n, $f_n(v)$ is any point within $f(v)$ (the f defined in the statement).

(b) Extend f_n linearly to the points in the relative interior of the faces.

Note that each f_n is a continuous function, as it is has been defined in a piecewise linear way. Further, from Brouwer's Theorem, let x_n be a fixed point of f_n.

If any of the x_n is a vertex of Σ_n, then we are done, because, by construction, for a vertex v of Σ_n, we have $f_n(v) \in f(v)$, and if v is fixed point, $v = f_n(v)$.

Thus we may assume that no x_n is a vertex of Σ_n, and thus belongs to a face of dimension d and is the convex combination of vertices, as follows

$$x_n = a_{n,0}x_{n,0} + ... + a_{n,d}x_{n,d}.$$

Applying f to this equation and using linearity, we have

$$x_n = a_{n,0}y_{n,0} + ... + a_{n,d}y_{n,d}.$$

Upon taking the limit as $n \to \infty$, we may assume that we have a subset of indices for which each sequence of x_ns, $a_{n,i}$s and $y_{n,i}$s converges, say to points x^*, a_i^* and y_i^*. Moreover, x^* is also the limit of the corresponding subsequences of $x_{n,i}$ for all i, and the maximum diameter of Σ_n tends to 0. To finish the argument, now we consider the facts:

- $x^* = a_0^*y_0^* + ... + a_d^*y_d^*$.
- $y_i^* = \lim_{n\to\infty} y_{n,i} = \lim_{n\to\infty} f_n(x_{n,i})$, and $f_n(x_{n,i}) \in f(x_n)$ by construction.
- $\lim_{n\to\infty} x_{n,i} = x^*$.

Condition (ii) of the statement now gives $y_i^* \in f(x^*)$ for all i. Because $f(x^*)$ is convex, it follows that $x^* \in f(x^*)$, as we wanted. □

12 Two-person zero sum games

Game theory is the mathematical formalization of the interaction among players within a set of decision making and payoff. Virtually any game, but more generally, any process in which individual strategies combine to generate a global setting that pays or charges something from each player can be formalized, and its properties studied, from the a mathematical point of view.

To start with a simple example, consider the following matrix, call it M, that represents rock, paper, scissors (with respect to player A, corresponding to the rows).

$$\begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

A *mixed strategy* is a probability distribution over the possible strategies. For example, player A might decide they will play rock or scissors, with probability 1/2 each, whereas player B decides to play scissors with probability 1/2, and each of the other possibilities with equal probability. It is immediate to verify that the expected payoff for player A corresponds to

$$\begin{pmatrix} 1/2 & 1/2 & 0 \end{pmatrix} \cdot M \cdot \begin{pmatrix} 1/4 \\ 1/2 \\ 1/4 \end{pmatrix} = 1/8.$$

More generally, if M is a $m \times n$, then $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^n$ is a pair of mixed strategies if $x \geq 0$ and $y \geq 0$, and $\sum x_i = \sum y_j = 1$ (I will denote these convex sets by C_m and C_n respectively). In this case, the expected payoff for A is

$$V(x, y) = x^T M y = \sum_{i=1}^{m} \sum_{j=1}^{n} M_{ij} x_i y_j.$$

It is immediate to verify that $-M^T$ is the payoff matrix for player B, which is the defining property for a two-person game to be called a zero-sum game (whenever a player wins quantity w, the other loses quantity w).

A pair of mixed strategies (x^*, y^*) is called an *equilibrium* point if

$$V(x, y^*) \leq V(x^*, y^*) \quad \text{for all} \quad x \in C_m$$

and

$$V(x^*, y^*) \leq V(x^*, y) \quad \text{for all} \quad y \in C_n.$$

Lemma 12.1. *Given a two-player zero sum game with expected payoff function V (which is continuous!), as defined above, an equilibrium pair exists if and only if*

$$\max_{x \in C_m} \min_{y \in C_n} V(x, y) = \min_{y \in C_n} \max_{x \in C_m} V(x, y).$$

Proof. If the equilibrium exists, then

$$\min_{y \in C_n} \max_{x \in C_m} V(x,y) \leq \max_{x \in C_m} V(x,y^*) = V(x^*,y^*) = \min_{y \in C_n} V(x^*,y) \leq \max_{x \in C_m} \min_{y \in C_n} V(x,y).$$

On the other hand, it is immediate that

$$\max_{x \in C_m} \min_{y \in C_n} V(x,y) \leq \min_{y \in C_n} \max_{x \in C_m} V(x,y).$$

Thus equality holds.

To see the converse, let x^* be the solution obtained in C_m for the maximin, and y^* be obtained in C_n for the minimax. We have

$$V(x^*,y^*) \geq \min_{y \in C_n} V(x^*,y) = \max_{x \in C_m} \min_{y \in C_n} V(x,y) = \min_{y \in C_n} \max_{x \in C_m} V(x,y) = \max_{x \in C_m} V(x,y^*) \geq V(x^*,y^*).$$

With equality holding, we have that $V(x^*,y^*) \leq V(x^*,y)$ for all i, and $V(x,y^*) \leq V(x^*,y^*)$ for all x, as we wanted. □

John von Neumann proved in 1928 a result stating that two-person zero sum games always admit an equilibrium — thus possibly founding game theory. Jonh Nash later showed how to prove his result using Brouwer's fixed point theorem.

Theorem 12.2 (von Neumann)**.** *Let M be an $m \times n$ matrix representing the payoff of player A in a two-person zero sum game. Then there is a pair of mixed strategies which gives an equilibrium to the game.*

Proof. Consider the compact convex set $C_m \times C_n$ of $\mathbb{R}^{m+n}$. Our goal is to define a function T from this set to itself so that a fixed point of this function corresponds precisely to an equilibrium of the game. The fact that this fixed point exists is an immediate consequence of Brouwer's fixed point theorem.

For $(x,y) \in C_m \times C_n$, let

$$c_i(x,y) = \max\{V(e_i,y) - V(x,y), 0\} \quad \text{and} \quad d_j(x,y) = \max\{V(x,y) - V(x,e_j), 0\}.$$

Define $T(x,y) = (x',y')$ with

$$x'_i = \frac{x_i + c_i(x,y)}{1 + \sum_{k=1}^m c_k(x,y)} \quad \text{and} \quad y'_j = \frac{y_j + d_j(x,y)}{1 + \sum_{k=1}^n d_k(x,y)}.$$

It is immediate to verify that T defined as such maps $C_m \times C_n$ to itself.

Moreover, if x^* and y^* are equilibrium points, then $c_i(x^*,y^*) = 0$ for all i, and $d_j(x^*,y^*) = 0$ for all j, thus $T(x,y) = (x,y)$.

It remains to show that a fixed point necessarily determines an equilibrium. Assume $(\overline{x},\overline{y})$ is a fixed point. Note that, because V is linear on x, then for at least one index with $\overline{x}_i > 0$ it must be $V(\overline{x},\overline{y}) \geq V(e_i,\overline{y})$. Thus for this index it follows that $c_i(\overline{x},\overline{y}) = 0$, and hence we can conclude that

$$0 \neq \overline{x}_i = \frac{\overline{x}_i}{1 + \sum_{k=1}^m c_k(\overline{x},\overline{y})}.$$

Therefore $\sum_{k=1}^m c_k(\overline{x},\overline{y}) = 0$, so $c_k(\overline{x},\overline{y}) = 0$ for all indices, showing that $\overline{x}$ maximizes $V(\cdot,\overline{y})$. Likewise for the other direction, thus $(\overline{x},\overline{y})$ is an equilibrium. □

13 Nash Equilibrium

In the previous section, we discussed a very specific type of games. In general, the payoffs might not be symmetric (if one player loses, it doesn't mean the other wins the same amount).

The standard example of an "all-players always lose" game is the prisoner's dilemma. Here, two prisoners are given a choice of either staying silent or telling on their accomplice. The numbers inside the parenthesis indicate their penalty for each combined stragy (assume prisoner A is the first coordinate).

$$\begin{array}{ccc} & & \text{B} \\ & \text{silence} & \text{tell on the other} \\ \text{A} \quad \text{silence} & (-2,-2) & (-6,-1) \\ \text{tell on the other} & (-1,-6) & (-4,-4) \end{array}$$

Clearly the best combined strategy is both staying silent, however, this not a strategy that attains equilibrium, because both have an incentive to change their behaviour.

Moreover, games can have more than 2 players. Our first goal now is to formalize the most general type of game with a finite number of players and strategies. Here we go.

(a) A set of players $P = \{1, ..., n\}$.

(b) A set of *strategies* to each player i, say $S_i = \{1, ..., k_i\}$.

(c) A *strategy profile* is a choice a strategy for each player, thus the set of all strategies profiles is $S = S_1 \otimes S_2 \otimes ... \otimes S_n$.

(d) A payoff function, that, for each strategy profile, determines how much each player makes or loses, that is, a function

$$p : S \to \mathbb{R}^n.$$

We assume that given a profile $s \in S$, player j makes $p_j(s)$.

Each player may chose a strategy, or a probability distribution over its possible strategies. If player i choses distribution $\sigma_i : S_i \to [0, 1]$, then the combined distribution on S is given by $\sigma : S \to [0, 1]$, with

$$\sigma(s) = \sigma(\ell_1, ..., \ell_n) = \prod_{i=1}^{n} \sigma_i(\ell_i),$$

for all $s = (\ell_1, ..., \ell_n) \in S = S_1 \otimes ... \otimes S_n$.

Given that each player i chose distribution σ_i, that is, with $\sigma = (\sigma_1, ..., \sigma_n)$, we have that the expected payoff to player j is given by

$$e(j; \sigma) = e(j; \sigma_1, ..., \sigma_n) = \sum_{s \in S} \sigma(s) p_j(s).$$

Let Σ_i be the set of all probability distributions over S_i, henceforth referred to the set of all possible strategies of player i. Note that Σ_i is a k_i-dimensional simplex (in particular, a compact and convex set).

If players 2 up to n choose strategies σ_2 up to σ_n, then player 1 would like to find $\overline{\sigma_1} \in \Sigma_1$ so that

$$e(1; \overline{\sigma_1}, \sigma_2, ..., \sigma_n) \geq e(1; \sigma', \sigma_2, ..., \sigma_n)$$

for all $\sigma' \in \Sigma_1$. This is called a *best response* of player 1 to strategies $(\sigma_2, ...\sigma_n)$. The set of all best responses of player 1 against these strategies is denoted by

$$\Delta_1(\sigma_2, \sigma_3, ..., \sigma_n).$$

It is immediate to verify that this is a convex set. Moreover, the expected payoff function is continuous for each coordinate, thus its maximum is attained within the set Σ_1 (which is compact), therefore $\Delta_1(\sigma_2, \sigma_3, ..., \sigma_n)$ is always non-empty. Similar observations hold for all others Δ_is.

A mixed strategy profile $(\overline{\sigma_1}, ..., \overline{\sigma_n})$ is a *Nash equilibrium* if, for all $i \in [n]$, we have

$$e(i; \overline{\sigma_1}, ..., \overline{\sigma_i}, ..., \overline{\sigma_n}) \geq e(i; \overline{\sigma_1}, ..., \sigma', ..., \overline{\sigma_n}) \quad \text{for all} \quad \sigma' \in \Sigma_i.$$

The obvious observation follows.

Proposition 13.1. *A mixed strategy profile $(\overline{\sigma_1}, ..., \overline{\sigma_n})$ is a Nash equilibrium if and only if, for all i,*

$$\overline{\sigma_i} \in \Delta_i(\overline{\sigma_1}, ..., \overline{\sigma_{i-1}}, \overline{\sigma_{i+1}}, ..., \overline{\sigma_n}).$$

□

We are now in a very convenient position to show John Nash's famous equilibrium result.

Theorem 13.2 (Nash 1950)**.** *A Nash equilibrium exists for all finite games.*

Proof. Let $\Sigma = \Sigma_1 \times ... \times \Sigma_n$ be the space of all mixed strategies. Consider the function $F : \Sigma \to 2^\Sigma$ defined by the rule

$$\begin{pmatrix} \sigma_1 \\ \sigma_2 \\ \vdots \\ \sigma_n \end{pmatrix} \mapsto \Delta_1(\sigma_2, \sigma_3, ..., \sigma_n) \times \Delta_2(\sigma_1, \sigma_3, ..., \sigma_n) \times ... \times \Delta_n(\sigma_1, \sigma_2, ..., \sigma_{n-1}).$$

From the previous proposition, we would like to show that there is $\sigma \in \Sigma$ so that $\sigma \in F(\sigma)$. So we only need to show that this function F, and the set Σ, satisfy the hypothesis of Kakutani's fixed point theorem.

First, note that Σ is compact and convex. Second, as we have already observed, each Δ_i is non-empty and convex. So it remains to show condition (ii) of Kakutani's theorem statement.

Assume (σ^n, τ^n) is a sequence of pairs of elements in Σ that converges to (σ, τ), but for the sake of a contraction, assume $\tau^n \in F(\sigma^n)$ for all n but $\tau \notin F(\sigma)$. Thus there is an index i so that $\tau_i \notin \Delta_i(\sigma_j)_{j \neq i}$, hence there is a better response to τ_i given choices σ_j for the players $j \neq i$, call it x_i. Because τ_i is the limit of the sequence $(\tau^n)_i$, it must be that x_i eventually becomes better than the elements of the sequence for some large n (as the expected payoff function is continuous), hence a contradiction to the fact that $\tau^n \in F(\sigma^n)$ for all n.

The hypothesis in Kakutani's fixed point theorem hold, therefore there is $\sigma \in \Sigma$ so that $\sigma \in F(\sigma)$, which from the previous proposition must be an equilibrium of the game. □

John Nash's original proof did not use Kakutani's fixed point theorem. In fact, it is possible to generalize the proof strategy of von Neumann's minimax theorem to prove Nash's equilibrium theorem.

Exercise 13.3. Prove Theorem 13.2 using Brouwer's fixed point theorem directly.

On the other hand, almost the same proof of Theorem 13.2 works to show the more general result. An Euclidean Game is defined as

(a) A set of players $P = \{1, ..., n\}$.

(b) A set of strategies to each player i, say $S_i \subseteq \mathbb{R}^{k_i}$.

(c) A payoff function, that, for each strategy profile in $S = S_1 \otimes S_2 \otimes ... \otimes S_n$, determines how much each player makes or loses, that is, a function

$$p : S \to \mathbb{R}^n.$$

We assume that given a profile $s \in S$, player j makes $p_j(s)$.

Theorem 13.4 (Debreu, Glicksberg, Fan (all independently) 1952)**.** *Given an Euclidean Game, if each S_i is non-empty, compact and convex, and each p_j is continuous and (quasi)concave, then the game admits a Nash equilibrium.* □

14 Gale and Shapley

Certain games are called cooperative games or coalitional games. In these, the players most form alliances and derive benefits from each, thus implying in the existence of a set of preferred alliances to each player. Naturally, constraints to the number of alliances a player can participate in exist, thus creating a game theoretic model of competition for good alliances.

A simple example to illustrate a coalitional games is that in which players are vertices of a bipartite graph G, and they can form alliances with at most one player to which it is adjacent in the graph. For each $v \in V(G)$, there exists a strict linear ordering $<_v$ on $E(G)$ which determines which are the preferred edges to player v.

One is of course interested in finding a choice of edges that achieves an equilibrium, that is, a choice in which no pair of players which are not matched would rather undo their current matchings to form a new one between them. That is, one is interested in finding a matching M so that, for all edges $e \notin M$, there is at least one edge $f \in M$ so that $f \cap e = \{v\}$ and v and $e <_v f$. Such a matching is called stable , and otherwise it is called unstable. Note that a stable matching is always maximal (though it not need to have maximum size).

An old result in graph (and game) theory states that stable matchings always exist for bipartite graphs.

Theorem 14.1 (Gale and Shapley 1962)**.** *Let G be a bipartite graph. For any set of preferences, G has a stable matching.* □

This theorem can be proved elementarily (and algorithmically). Let $A, B \subseteq V(G)$ be the bipartition. Vertex $a \in A$ is acceptable to a vertex $b \in B$ if $ab \in E \backslash M$ and b prefers ab over any edge currently at it in M. The basic idea is to start with the empty matching and add edges greedily as follows. Find an unmatched vertex $a \in A$ for the current matching M which is acceptable to at least one vertex in B (if no such vertex exists, terminate). Add to M the maximal edge according to $<_a$, say edge ab, so that a is acceptable to b, and discard any edge of M currently at b.

It is a worthy exercise to show that this procedure indeed terminates with a stable matching for G.

Our goal, however, is to analyse cooperative games in a more general framework, and to prove an equilibrium result using topological methods.

15 Cooperative games, the core and Scarf's lemma

We start with a definition that generalizes the scenario described in the context of Gale & Shapley. A *finitely generated non-transferable utility game* is given by:

- A set of n players $\{1, ..., n\}$.
- A multiset of coalitions $S_j \subseteq \{1, ..., n\}$, $j = 1, ..., m$. We can interpret a coalition as an action performed by the players which form it.
- Each player has a total ordering $<_i$ of the coalitions that they participate in, meaning that $S_j <_i S_k$ indicates player i prefers coalition S_k to S_j (implicit in this is that $i \in S_j$ and $i \in S_k$).

A set of coalitions $\mathcal{S}$ is in the *core of the game* if they are disjoint and for each $S' \notin \mathcal{S}$, there is a player $i \in S'$ and a coalition $S \in \mathcal{S}$ so that $S' <_i S$.

As we have seen above, Gale & Shapley imply that if coalitions are edges of bipartite graph, then the core is always non-empty. The result of course does not generalize immediately to any finitely generated cooperative game, but we can fix it introducing the notion of a fractional core.

A vector $x : \{1, ..., m\} \to \mathbb{R}_+$ is in the *fractional core* if, for each player $i \in \{1, ..., n\}$, we have

$$\sum_{j:i \in S_j} x_j \leq 1,$$

and, for each j, there is a player $i \in S_j$ with

$$\sum_{k:i \in S_k} x_k = 1$$

and $S_j \leq_i S_{k'}$ whenever $i \in S_{k'}$ and $x_{k'} > 0$.

The motivation is that coalitions now have intensities, players have capacities, and there is no coalition in which every player wants to increase its intensity. Note that integer valued elements in the fractional core are elements of the core. Our first goal is to show the following result.

Theorem 15.1. *The fractional core of a finitely generated non-transferable utility game is always non-empty.*

In order to prove this (and other results), we will make use of a lemma due to Scarf.

Let C be a matrix. A set S of columns is called a *dominating set of columns* if, for every column index j, there exists a row index i such that $c_{ij} \leq c_{ik}$ for all k in S.

Note that the collection of all dominating sets of columns in C is non-empty, as the single column that contains the largest element of a row is a dominating set. Also, this collection forms an abstract simplicial complex.

Finally, if in each row all elements are distinct, then the maximum size of a dominating set of columns is at most the number of rows of the matrix C.

Lemma 15.2 (Scarf)**.** *Let B and C be $n \times (n+m)$ matrices, with entries in $\mathbb{R}$, and let $b \in \mathbb{R}^n$, so that*

- *$b > 0$ (meaning, all entries of b are positive),*
- *the first n columns of B form an identity matrix,*
- *the set of non-negative $x \in \mathbb{R}^{n+m}$ satisfying $Bx = b$ is bounded,*
- *all entries in each row of C are distinct,*
- *in each row i of C, c_{ii} is the smallest element, and the $(n-1)$ largest elements are $\{c_{i1}, ..., c_{in}\} \backslash \{c_{ii}\}$.*

Then, there exists a non-negative x so that $Bx = b$ and $\operatorname{supp}(x)$ indexes a dominating set of columns in C.

Exercise 15.3. Explain why the first n columns of C is not a dominating set of columns. What would be a dominating set of columns containing the first $n-1$ columns of C?

Proof of Theorem 15.1. Let B be a $n \times (n+m)$ matrix whose rows are indexed by players, and last m columns are indexed by coalitions. Make the first n columns of B equal to the identity matrix, and have entry $B_{i,n+j} = 1$ if $i \in S_j$, and $= 0$ otherwise. Define matrix C', $n \times m$, so that rows are indexed by players, columns by coalitions, entries in each row are all different, and so that

- $C'_{i,j} < C'_{i,k}$ if $i \in S_j \cap S_k$ and $S_j <_i S_k$,
- $C'_{i,j} < C'_{i,k}$ if $i \in S_j$ but $i \notin S_k$.

Append to C' a square matrix to its left to form matrix C so that C satisfies that hypotheses in Scarf's lemma. (Short exercise: find an explicit definition for C that works.)

Make vector $b = \mathbb{1}$. Note that the set of y with $By = b$ is bounded (why?), so we can apply Scarf's lemma to recover a non-negative vector x with $Bx = b$ and so that $\operatorname{supp}(x)$ indexes a dominating set in C. Let x' be its restriction to columns indexed by players and z be its restriction to the columns indexed by coalitions.

It is immediate to verify that z is in the fractional core:

- $Bx = \mathbb{1}$ implies $\sum_{j:i \in S_j} z_j \leq 1$.
- Now fix coalition S_j. Then, there is player i so that $C'_{i,j} \leq C'_{i,k}$ for all k in the support of z. Thus, the ith entry in x' is 0, and thus $\sum_{k:i \in S_k} z_k = 1$. Moreover, i must be S_j, because $\operatorname{supp}(z)$ contains at least one coalition containing i, say S_k, and $C'_{i,k} \geq C'_{i,j}$. Because all entries in the rows of C are different, it follows that if $i \in S_j \cap S_k$, that is, if $z_k > 0$, then $C'_{i,j} < C'_{i,k}$ implies $S_j <_i S_k$, as we wanted.

□

16 Unique add and unique remove*

Our goal in this section is to derive a common framework over which Sperner's lemma and Scarf's lemma can be proved.

An r-uniform hypergraph $\mathcal{H}$ has the *unique add* property if, for each $e \in \mathcal{H}$, and each $v \in V(\mathcal{H})\backslash e$, there exists a unique $w \in e$ so that $e\backslash\{w\} \cup \{v\} \in \mathcal{H}$. Note that if a hypergraph is formed by taking all transversals of a fixed partition of a finite set, then this hypergraph has the unique add property.

An r-uniform hypergraph $\mathcal{H}$ has the *unique remove* property if, for each $e \in \mathcal{H}$, and each $w \in e$, there exists a unique $v \in V(\mathcal{H})\backslash e$ so that $e\backslash\{w\} \cup \{v\} \in \mathcal{H}$. For example, consider a triangulation Σ of the d-dimensional simplex σ with no points in $\partial\sigma$ other than its vertices $\{v_0, ..., v_d\}$. Let $\mathcal{H}$ be $\{\text{vts}(\tau) : \tau \in \Sigma, \dim \tau = d\} \cup \{v_0, ..., v_d\}$. Then $\mathcal{H}$ has the unique remove property.

Theorem 16.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be r-uniform hypergraphs on the same vertex set V. Suppose $\mathcal{A}$ has the unique remove property, and $\mathcal{B}$ has the unique add property, Then $|\mathcal{A} \cap \mathcal{B}|$ is even.*

Proof. Define a directed graph D as follows.

- Vertex set is $\mathcal{A} \cup \mathcal{B}$.
- Fix $v_0 \in V$. Put an arc $A \to B$ between two vertices of D if there exists $w \in V$, $w \neq v_0$, such that $B = A\backslash\{w\} \cup \{v_0\}$.

Note that if $v_0 \in A$, then the outdegree of A is 0. If $v_0 \notin A$, then there are two cases. If $A \cup \{v_0\}$ does not contain any $B \in \mathcal{B}$, then the outdegree of A is 0. If it contains some $B \in \mathcal{B}$, then let $w = A \cup \{v_0\}\backslash B$. Due to the unique add property, there is a unique $v \in B$ so that $B\backslash\{v\} \cup \{w\} = B' \in \mathcal{B}$. If $v \neq v_0$, then B' is the unique other out neighbour of A, and so the outdegree of A is 2. If $v = v_0$, then $B' = A \in \mathcal{B}$, and the outdegree of A is equal to 1 (and also conversely). By an analogous analysis, if $B \in \mathcal{B}$, then its indegree is either 0 or 2, or

$$d^-(B) = 1 \iff B \in \mathcal{A} \text{ and } v_0 \in B.$$

If the directions are removed and the graph is made undirected, then the only vertices of odd degree are those in $\mathcal{A} \cap \mathcal{B}$, thus this has even cardinality. □

Note that replacing "unique" by "odd" in the unique add or remove properties would have yield the same result.

With this theorem, we can provide a new proof of Sperner's lemma. Let Σ be a triangulation of a geometric d-dimensional simplex σ in $\mathbb{R}^d$ with vertex set $V(\sigma) = \{v_0, ..., v_d\}$. Assume Σ has no other vertex in $\partial\sigma$. Let c be a colouring of the points of Σ with $r = d + 1$ colours so that the vertices of σ all get different colours. Define $\mathcal{A}$ and $\mathcal{B}$ r-uniform hypergraphs on $V(\Sigma)$, with $\mathcal{A}$ having all d-dimensional simplices (and $V(\sigma)$), and $\mathcal{B}$ having all multicoloured r-sets of $V(\Sigma)$. By Theorem 16.1, as $\{v_0, ..., v_d\}$ is already multicoloured, we can guarantee the existence of a multicoloured simplex within Σ.

Moreover, to show the general version of Sperner's lemma, simply start with Σ and create σ, and create d-dimensional simplices picking vertices from σ and lower dimensional simplices from $\partial\mathcal{A}_{[m]}$, as the picture below shows:

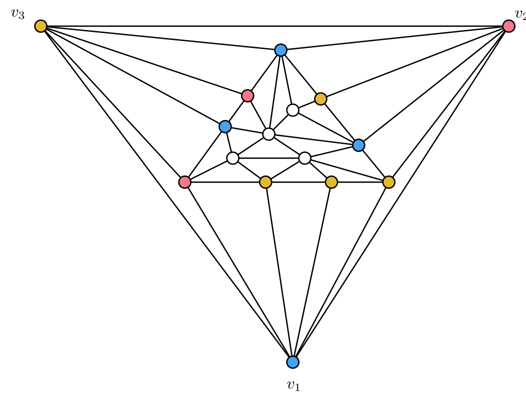

For example, in this picture, there are an odd number of multicoloured simplices on each side of the original triangle using only one new vertex, thus an odd number in total (as there are three faces), and also an odd number using two of the new vertices. Finally, simplex $\{v_1, v_2, v_3\}$ is multicoloured, resulting in an odd number in total. Thus there must be an odd number of multicoloured simplices within the triangulation.

Recall now the statement of Scarf's lemma, with the matrices B and C. We will show Scarf's lemma making use of the following two lemmas.

Lemma 16.2. *Let $\mathcal{C}$ the set of all n-subsets of columns of C that are dominating, along with the set that contains the first n columns. Then $\mathcal{C}$ has the unique remove property.*

Lemma 16.3. *Let $\mathcal{B}^*$ denote the set of all n-subsets S of columns of B such that S is linearly independent, and S contains the support* $\text{supp}(x)$ *of some non-negative solution x to $Bx = b$. Then $\mathcal{B}^*$ contains an n-uniform hypergraph $\mathcal{B}$ that has the unique add property, and, moreover, any element of $\mathcal{B}^*$ that is* $\text{supp}(x)$ *for some non-negative x with $Bx = b$ is also in $\mathcal{B}$.*

Proof of Scarf's lemma 15.2. Seeing $\mathcal{B}$ and $\mathcal{C}$ as n-uniform hypergraph over the same index set of both matrices B and C, Theorem 16.1 and both lemmas above show that $|\mathcal{B} \cap \mathcal{C}|$ is even. But the set of the first n columns belong to this intersection (as $b > 0$).

So there exists S that is dominating in C and contains $\text{supp}(x)$ for some non-negative x solution to $Bx = b$. Thus $\text{supp}(x)$ is also dominating in C. □

So we are left with the task of showing Lemmas 16.2 and 16.3.

17 Proofs of the lemmas*

Proof of Lemma 16.2. Let σ_0 be the set of the first n columns. Let $A \in \mathcal{C}$ and $y \in A$. We split into cases.

1. $A = \sigma_0$. Let t be the column $> n$ that contains the largest entry of row y. Then $A - y + t$ is dominating, thus in $\mathcal{C}$, and $A - y + s$ does not dominate t for any other s.

2. $A \neq \sigma_0$. For each $z \in A$, there is a unique row for which z is dominated by A in that row (because entries are different in all rows and there are n columns in A). Let i be the row in which y is dominated by A, thus, the ith entry of y is the smallest amongst the other columns in A. Let u be the column in A whose i-entry is the second smallest, and let j be the row in which u is dominated by A. Note that u is dominated by $A \backslash \{y\}$ in both rows i and j, and no other column of A is dominated in row j.

   (a) $j \in A$. Since j dominates u in row j, it must be that $j = u$. But u had the second smallest entry in row i among all columns of A. So A actually consists of $\sigma_0 - i + y$. Hence column i is such that $A - y + i = \sigma_0$. Moreover $A - y + s$ does not dominate column y for any $s > n$.

   (b) $j \notin A$. Let W denote the set of columns $w \notin A$ that are dominated by $A - y$ only in row j. Note that if $j \notin W$, then $A - y \subseteq \{1, ..., n\}$ (again), so we are to σ_0. So we may assume $j \in W$. Let $t \in W$ be the column in W with the largest entry in row j. We claim $A - y + t \in \mathcal{C}$, and is unique. Note that $y \notin W$ since $y \in A$. We will show $A - y + t$ is dominating. Let v be a column of C. Then v is dominated by $A - y$:

   - if v is dominated by $A - y$ in some row $r \neq j$, then since $t \in W$, t is not dominated by $A - y$ in row r, so $A - y + t$ dominates v in row r.
   - if v is dominated by $A - y$ only in row j, then note that $v \in W$ since $v \notin A$ (we observed that no element of A was dominated by $A - y$ only in row j). Therefore t dominates v in row j,

   To see uniqueness, suppose $z \notin A$ and $A - y + z$ is dominating. Then z is dominated by $A - y + z$ in row i or in row j, but only one of them. Suppose z is dominated in row i. Then $A - y + z$ dominates y and this can be only in row i. So ith entry of z is $>$ than the ith entry of y. But A used to dominate z, and hence that could only be in row i, a contradiction. So z is dominated in row j. Therefore $z \in W$, and hence $z = t$ by a similar argument. So the only thing remaining to check is that $A - y$ cannot become σ_0 in this case. In fact, if $A - y \subseteq \{1, ..., n\}$, then $A = \{1, ..., n\} \backslash \{j\} \cup \{t_j\}$, where t_j is the column $> n$ with the largest entry in row j. But $i \in A$ and y is dominated in row i, which is not possible, since c_{ii} is the smallest entry in row i.

□

Proof of Lemma 16.3. First we prove a special case (1), and then we show how to reduce the general case (2) to it.

We say (B, b) is a nondegenerate system if every set of n linearly independent columns of B is the support of a solution y to $By = b$.

(1) Let B be an $n \times (n+m)$ real matrix, $b \in \mathbb{R}^n$ with $b > 0$, and assume (B, b) is nondegenerate. Assume the set of nonnegative solutions x to $Bx = b$ is bounded. Define the n-uniform hypergraph $\mathcal{H}$ on the columns of B to have hyperedges corresponding to sets of linearly independent columns which are the support of a nonnegative solution x to $Bx = b$. Then $\mathcal{H}$ has the unique add property.

To see that, let $S \in \mathcal{H}$ and let u be a column of B with $u \notin S$. Let x be so that $\text{supp}(x) = S$, $x \geq 0$, and $Bx = b$. First we prove that there is at least one element $v \in S$ such that $S - v \cup u$ is linearly independent and is the support of a non-negative vector y such that $By = b$. Let S' be a subset of S of minimal size such that $S' \cup u$ is linearly dependent.

(a) Let z be a vector with support equal to $S' \cup u$ and $Bz = 0$. Prove that z has positive and negative entries.

(b) Assume wlog $z_u > 0$. Choose $vinS'$ so that $z_v < 0$ and $x_v/|z_v|$ is minimum. Define $y = x + (x_v/|z_v|)z$. Prove that $By = b$.

(c) Prove that $y \geq 0$.

(d) Prove that $S - v \cup u$ is linearly independent.

(e) Using that (B, b) is nondegenerate, prove that $\text{supp}(y) = S - v \cup u$.

(f) Assume now that there is $v' \in S$, $v' \neq v$, so that $S - v' \cup u$ is the support of a nonnegative vector y' with $By' = b$. Find a contradiction (use different linear combinations of x, y, y', and the fact that S is a linear independent set.)

Note that you have just finished proving (1).

Now on to the second step. Let B be an $n \times (n+m)$ real matrix, $b \in \mathbb{R}^n$ with $b > 0$. Assume the set of nonnegative solutions x to $Bx = b$ is bounded. Define the n-uniform hypergraph $\mathcal{H}$ on the columns of B to have hyperedges corresponding to sets of linearly independent columns which contain the support of a nonnegative solution x to $Bx = b$. In principle, we cannot show directly that B has the unique add property without the hypothesis that (B, b) is nondegenerate.

(2) If ϵ is a vector, let $b_\epsilon = b + \epsilon$. Then there exists ϵ (with small entries) so that the hypergraph $\mathcal{H}'$ which consists of the linearly independent sets of n columns of B which are the support of a nonnegative solution x to $Bx = b_\epsilon$ has the following properties

(i) $\mathcal{H}' \subseteq \mathcal{H}$.

(ii) if $S \in \mathcal{H}$ and $S = \text{supp}(x)$ for some nonnegative solution x to $Bx = b$, then $S \in \mathcal{H}'$.

(iii) (B, b_ϵ) is nondegenerate (thus $\mathcal{H}'$ has the unique add property.)

A formal proof of this result requires some basics of topology in $\mathbb{R}^n$, in particular, the fact the linear transformations are continuous functions, and that if a continuous functions maps a set X to a set which is open (closed), then X itself is open (closed). Anyway, I will allow you to give not so formal arguments if you are not comfortable with these concepts.

(a) Prove that there is an open ball W centred at b so that if $Bx \in W$ for some x whose support is contained in a set S of linearly independent columns of B, then $S \in \mathcal{H}$.

(b) Let S be a set of linearly independent columns of B which is the support of a nonnegative solution y to $By = b$. Let $\delta > 0$ be smaller than the smallest positive entry of y. Let Y_S be obtained upon perturbing possibly all entries of y by adding or subtracting δ. Prove that there is an open ball Y centred at b which is contained in Y_S for all S.

(c) Prove that there is an open ball Z centred at b which is contained in W and in Y, and, moreover, show that there is at least one $\epsilon \in Z$ so that b_ϵ can only be written as a linear combination of n linearly independent columns of B.

(d) Finish the argument to prove (2).

□

Exercise 17.1. Prove (a)-(f) in (1) and (a)-(d) in (2) above.

18 Combinatorial applications of Scarf's Lemma*

Let D be a directed graph. A *kernel* K in D is an independent absorbing set:

- no arc joins two vertices in K, and
- for all $v \notin K$, there is an arc vw in the graph, with $w \in K$.

A *fractional kernel* is a function $\phi : V(D) \to [0,1]$ that is fractionally independent and fractionally absorbing, that is

- for each clique W of D (set of vertices so that between any two vertices in the set there is an arc), we have $\sum_{w \in W} \phi(w) \leq 1$, and
- for each $v \in V(D)$, $\sum_{w \in \Gamma^+(v)} \phi(w) \geq 1$, where $\Gamma^+(v) = \{v\} \cup \{w : vw \text{ is an arc}\}$.

Our goal is to show the existence of fractional kernels in certain directed graphs, but, in fact, we will see a stronger version, that is, we will show that for each $v \in V(D)$, there exists a clique $W \subseteq \Gamma^+(v)$ with $\sum_{w \in W} \phi(w) \geq 1$.

Exercise 18.1. Show that the cyclic orientation of K_3 admits no fractional kernel.

We say an orientation of a clique is transitive if there is an ordering of the vertices such that each arc obeys the ordering.

Theorem 18.2. *Let D be a directed graph in which all cliques are transitive. Then D has a strong fractional kernel.*

Proof. Construct matrices B and C as follows. Let W be the set of all maximal cliques in D. Then B and C will have W as row index set. The right hand side submatrix of B will be the clique vs vertex incidence matrix. Make $b = \mathbb{1}$. Note that if $(\psi \mid \phi)$ is a nonnegative solution to $Bx = b$, then, for each clique W, we have

$$\psi(W) + \sum_{w \in W} \phi(w) = 1,$$

thus ϕ is fractionally independent. Also, note easily that $\{x : Bx = b\}$ is bounded, as everything is nonnegative positive and there are no 0 columns.

Define the matrix C as follows. For row W and column $v \in V$, if $v \in W$, then (W, v) entry is the rank of v in W (arcs go from smaller to higher ranks). If $v \notin W$, the (W, v) entries are an arbitrary permutation of $\{|W| + 1, ..., |V(D)|\}$. The (W, W) entry is 0, and the big entries in the left hand side block are all distinct and bigger than $|V(D)|$.

By Scarf's Lemma, there exists a nonnegative solution $(\psi \mid \phi)$ to $Bx = b$ such that $S = \operatorname{supp}(\psi \mid \phi)$ is dominating in C. By definition of B, ϕ is fractionally independent, so we must now verify that ϕ is strong fractionally absorbing.

Let $v \in V$. Then there is a row W in C such that S dominates v in W.

- If $v \in W$, the W-entry of each column of S is $\geq$ than the W-entry of v, so all $w \in W$ with $\phi(w) > 0$ have higher rank than v (or $= v$), and so, for them, vw is an arc. Note that column W is not in S, since (W, W) entry is 0, so $\psi(W) = 0$, hence

$$\sum_{w \in W} \phi(w) = \sum_{w \in W \cap S} \phi(w) = 1.$$

Then, the set $W \cap S$ satisfies the condition for v (it forms a clique in $\Gamma^+(v)$).

- If $v \notin W$, by definition of C, no column of S can be in W. So $\sum_{w \in W} \phi(w) = 0$, thus $\psi(W) = 1$, so $W \in S$. But this is a contradiction, since column W cannot dominate v in row W.

So ϕ must be a strong fractional kernel. □

We present another application. Let V be a ground set, and $\leq_1$ and $\leq_2$ be two partial orders on V. A *dominating common antichain* of $\leq_1$ and $\leq_2$ is a subset A of V that is an antichain for both partial orders, and is so that for any $v \in V$, there is $a \in A$ with $v \leq_1 a$ or $v \leq_2 a$. The following result was proved by Sands (essentially), and it is easy to see that it generalizes the Gale & Shapley result.

Theorem 18.3. *For any two partial orders $\leq_1$ and $\leq_2$ on the same finite ground set V, there exists a dominating antichain for them.*

The theorem does not generalize for more than two orders directly, but its fractional version does. Given partial orders $\{\leq_j\}_{j=1}^k$ on V, a vector $x \in \mathbb{R}_+$ is a fractional dominating antichain if x is a fractional antichain for all $\leq_j$, that is, if

$$\sum_{c \in C} x_c \leq 1$$

for all chains C of all $\leq_j$, and if x is a fractional upper bound for any element of V, that is, if for all v, there is a chain $v = v_0 \leq_j v_1 \leq_j \ldots \leq_j v_\ell$ of some $\leq_j$ with

$$\sum_{i=0}^{\ell} x_{v_i} = 1.$$

Theorem 18.4. *Given a finite ground set V, and finite set $\{\leq_j\}_{j=1}^k$ of partial orders on V admits a fractional dominating antichain.*

Exercise 18.5. Prove this.

19 Borsuk-Ulam

At any given moment, two antipodal points on Earth's Equator have the same temperature. This sentence does not seem to be much surprising, however, it can be extended to a less believable claim: at any given moment, two antipodal points on Earth's surface have the same temperatura and air pressure.

In fact, these statement are anecdotes about the well known theorem due to Borsuk-Ulam. Recall that S_n is the sphere in $\mathbb{R}^{n+1}$, and let B_n denote the ball in $\mathbb{R}^n$.

Theorem 19.1. *The following statement are all true (and also easily shown to be equivalent to one another in terms of difficulty in proving).*

(a) For every continuous mapping $f : S_n \to \mathbb{R}^n$, there exists a point $x \in S_n$ with $f(x) = f(-x)$.

(b) For every antipodal mapping $f : S_n \to \mathbb{R}^n$, that is, f is continuous and $f(-x) = -f(x)$ for all $x \in S_n$, there exists a point $x \in S_n$ with $f(x) = 0$.

(c) There is no antipodal mapping $f : S_n \to S_{n-1}$.

(d) There is no continuous mapping $f : B_n \to S_{n-1}$ that is antipodal on the boundary, that is, satisfies $f(-x) = -f(x)$ for all $x \in S_{n-1} = \partial B_n$.

(e) For any cover $F_1, ..., F_{n+1}$ of the sphere S_n by $n + 1$ closed sets, there is at least one set containing a pair of antipodal points (that is, $F_i \cap -F_i \neq \emptyset$ for some i).

(f) For any cover $U_1, ..., U_{n+1}$ of the sphere S_n by $n + 1$ open sets, there is at least one set containing a pair of antipodal points.

First we will briefly comment on how these statements are all equivalent.

(a) $\iff$ (b)

Simply take $f(x)$ from (b), apply (a) to $f(x)$, and note that if $f(z) = -f(-z) = -f(z)$ for some z, then we have $f(z) = 0$. For the other direction, take $f(x)$ from (a), define $g(x) = f(x) - f(-x)$, and apply (b) to $g(x)$.

(b) $\iff$ (c)

One direction is obvious, as a mapping defined as in (c) is a mapping to $\mathbb{R}^n$ that is nowhere 0. For the converse, assume (b) fails for a function f, and define $g(x) = f(x)/\|f(x)\|$, which then will fail (c).

(c) $\iff$ (d)

Take the mapping $\pi : S_n \to B_n$ given by $(x_1, ..., x_{n+1}) \mapsto (x_1, ..., x_n)$. If (c) fails, say for a mapping f, then consider the function $g : B_n \to S_{n-1}$ given by $g(x) = f(\pi^{-1}(x))$. It is immediate to check that it fails (d). For the converse, if (d) fails with a function f, define $g : S_n \to S_{n-1}$ by $g(x) = f(\pi(x))$, and $g(-x) = -f(\pi(x))$. It is not difficult to check that it fails (c).

(a) $\implies$ (e)

Consider the cover as in (e), and define $f : S_n \to \mathbb{R}^n$ by $f(x) = (\text{dist}(x, F_1), ..., \text{dist}(x, F_n)$. From (a), let x be so that $f(x) = f(-x) = y$. If any coordinate of y is 0, then x is in the corresponding closed set, otherwise x and $-x$ lie in F_{n+1}.

(e) $\implies$ (c)

It is not difficult to see that there exists a covering of S_{n-1} by closed sets $F_1, ..., F_{n+1}$ so that no F_i contains a pair of antipodal points (example, project out the facets of n-simplex onto S_{n-1}). Thus, if (c) fails for a function f, the sets $f^{-1}(F_1), ..., f^{-1}(F_{n+1})$ would contradict (e).

(e) $\iff$ (f)

This is an exercise to those of you that like topology of $\mathbb{R}^n$.

Also, it is quite interesting to see a proof a Brouwer's fixed point theorem from Borsuk-Ulam.

Proof of Brouwer's fixed point theorem. Say $f : B_n \to B_n$ is continuous, without a fixed point. Let $g : B_n \to S_{n-1}$ be such that $g(x)$ is the point in which the ray originating in $f(x)$ and going through x intersects S_{n-1}. This $g(x)$ contradicts formulation (d) of Theorem 19.1 (note that $g(x) = x$ for all x on S_{n-1}, and g is continuous.) □

Our goal, starting from next section, is to prove Borsuk-Ulam Theorem by using a combinatorial lemma, in an analogous way to what we did with Sperner's Lemma and Brouwer's Theorem.

20 Tucker's Lemma*

A *triangulation* $\mathcal{T}$ of S_{n-1} is called antipodally symmetric if, for every $\sigma \in \mathcal{T}$, $-\sigma$ is also in $\mathcal{T}$.

Lemma 20.1 (Tucker)**.** *Let $\mathcal{T}$ be a triangulation of B_n such that the restriction of $\mathcal{T}$ to $\partial B_n = S_{n-1}$ is an antipodally symmetric triangulation. Let λ be a labelling of $V(\mathcal{T})$ with elements $\{\pm 1, ..., \pm n\}$, such that $\lambda(-v) = -\lambda(v)$ for every $v \in S_{n-1}$ (in other words, λ is antipodal on the boundary.) Then there exists a* 1*-simplex whose vertices are labelled i and $-i$ by λ, for some i.*

Let us now start the preparation to prove Tucker's lemma. Let $\mathcal{A}$ and $\mathcal{B}$ be simplicial complexes. A function $f : V(\mathcal{A}) \to V(\mathcal{B})$ is a *simplicial map* if $f(A) = \{f(v) : v \in A\}$ is a simplex of $\mathcal{B}$, for all $A \in \mathcal{A}$.

The n-dimensional *crosspolytope* is the convex hull of $\{\pm e_1, ..., \pm e_n\}$, where $e_1, ..., e_n$ is the canonical basis of $\mathbb{R}^n$. The boundary of the crosspolytope is denoted by $\diamond_{n-1}$. Note that $\diamond_{n-1}$ has a natural triangulation. Note that this triangulation can be viewed as an abstract simplicial complex with vertex $\{\pm 1, ..., \pm n\}$, where σ is a simplex if and only if *sigma* does not contain both i and $-i$, for all i.

Note finally that $\diamond_{n-1}$ is homeomorphic to the sphere S_{n-1}.

Lemma 20.2 (Tucker's lemma reformulated)**.** *Let $\mathcal{T}$ be a triangulation of B_n that is antipodally symmetric on the boundary. Then there does not exists a map $\lambda : V(\mathcal{T}) \to V(\diamond_{n-1})$ that is both a simplicial map and antipodal on the boundary.*

21 Chains*

Let $\mathcal{A}$ be a simplicial complex. A *k-chain* in $\mathcal{A}$ is a set C_k of k-dimensional simplices in $\mathcal{A}$.

- the empty chain is denoted by 0.
- for k-chains C_k and D_k, the symmetric difference of C_k and D_k is denoted by $C_k + D_k$.
- for $A \in \mathcal{A}$ of dimension k, the $(k-1)$-chain of all $(k-1)$-dimensional faces of A is denoted by ∂A.
- for a k-chain $\{A_1, ..., A_m\}$, its boundary is defined as $\partial A_1 + ... + \partial A_m$.
- for two k-chains C_k and D_k, note that $\partial(C_k + D_k) = \partial C_k + \partial D_k$.
- for a k-chain C_k, note that $\partial\partial C_k = 0$.

Suppose f is a simplicial map from $\mathcal{A}$ to $\mathcal{B}$. Fix k. Then f induces a map $f_{\#k}$ from the set of k-chains in $\mathcal{A}$ to the set of k-chains in $\mathcal{B}$ as follows:

- for a k-simplex $A \in \mathcal{A}$, and the k-chain $\{A\}$, define $f_{\#k}(\{A\}) = \{f(A)\}$ if $f(A)$ is a k-dimensional simplex of $\mathcal{B}$, and $= 0$ otherwise.
- for $C_k = \{A_1, ..., A_m\}$, define
$$f_{\#k}(C_k) = f_{\#k}(\{A_1\}) + ... + f_{\#k}(\{A_m\}).$$
- note that $f_{\#k}(C_k + D_k) = f_{\#k}(C_k) + f_{\#k}(D_k)$.

Lemma 21.1. *Let $\mathcal{A}$ be a simplicial complex. Let C_k be a k-chain in $\mathcal{A}$, and let f be a simplicial map from $\mathcal{A}$ to $\mathcal{B}$. Then*
$$f_{\#k-1}(\partial C_k) = \partial f_{\#k}(C_k)$$

Proof. Let A be a k-simplex of $\mathcal{A}$, consider $C_k = \{A\}$. Let $A = \{v_0, ..., v_k\}$.

- if $f(v_0)$, ... , $f(v_k)$ are all distinct, the $f_{\#k}(\{A\})$ is equal to $\{f(v_0), ..., f(v_k)\} \in \mathcal{B}$. So $\partial f_{\#k}(C_k)$ is the set of all k-subsets of $f(A)$. But ∂C_k is the set of all k-subsets of A, thus $f_{\#k-1}(\partial C_k)$ is the set of all k-subsets $f(A)$, since $f(v_0)$, ... , $f(v_k)$ are all distinct.
- if, say, $f(v_0) = f(v_1)$, then $f_{\#k}(C_k) = 0$. On the other hand, for a k-subset A' of A, if $v_0, v_1 \in A'$, then by definition $f_{\#k-1}(\{A'\}) = 0$. Thus
$$f_{\#k-1}\partial C_k = f_{\#k-1}(\{A \backslash v_0\}) + f_{\#k-1}(\{A \backslash v_1\}) = 0,$$
since $f(v_0) = f(v_1)$ and thus $f(A \backslash v_0) = f(\backslash v_1)$.

In general, for $C_k = \{A_1, ..., A_m\}$, we have

$$\begin{aligned}\partial f_{\#k}(C_k) &= \partial\big(f_{\#k}(\{A_1\}) + ... + f_{\#k}(\{A_m\})\big)\\ &= \partial f_{\#k}(\{A_1\}) + ... + \partial f_{\#k}(\{A_m\})\\ &= f_{\#k-1}\partial(\{A_1\}) + ... + f_{\#k-1}\partial(\{A_m\})\\ &= f_{\#k-1}\big(\partial(\{A_1\}) + ... + \partial(\{A_m\})\big)\\ &= f_{\#k-1}\partial C_k.\end{aligned}$$

□

Let $\mathcal{K}$ be a triangulation of S_{n-1} (formally, you can see $\mathcal{K}$ as the restriction of a triangulation of B_n, which is a "simplex", to its boundary). Thus, all $(n-2)$-dimensional simplex in $\mathcal{K}$ are contained in precisely two $(n-1)$-dimensional simplices, in particular, if A_{n-1} is the $(n-1)$-chain consisting of all $(n-1)$-dimensional simplices in $\mathcal{K}$, we have $\partial A_{n-1} = 0$.

Lemma 21.2. *Let $\mathcal{K}$ and $\mathcal{L}$ be triangulations of S_{n-1}. Suppose $f : V(\mathcal{K}) \to V(\mathcal{L})$ is a simplicial map. Let A_{n-1} be the $(n-1)$-chain consisting of all $(n-1)$-dimensional simplices in $\mathcal{K}$. Let $C_{n-1} = f_{\#n-1}(A_{n-1})$. Then either C_{n-1} is empty, or C_{n-1} consists of all the $(n-1)$-dimensional simplices in $\mathcal{L}$.*

Proof. Suppose C_{n-1} is non-empty. Assume that it does not contain all $(n-1)$-dimensional simplices in $\mathcal{L}$. As every $(n-2)$-dimensional simplex in $\mathcal{L}$ is contained in precisely two $(n-1)$-dimensional simplices, if C_{n-1} does not contain all $(n-1)$-dimensional simplices, then there must exist σ and σ', $(n-1)$-simplices in $\mathcal{L}$, that share a $(n-2)$-dimensional face τ, with $\sigma \in C_{n-1}$ and $\sigma' \notin C_{n-1}$. Then $\tau \in \partial C_{n-1}$, but

$$\partial C_{n-1} = \partial f_{\#n-1}(A_{n-1}) = f_{\#n-2}\partial(A_{n-1}) = f_{\#n-2}0 = 0,$$

a contradiction. □

If C_{n-1} is empty, as considered above, we say that f has even degree. If C_{n-1} consists of all the $(n-1)$-dimensional simplices in $\mathcal{L}$, we say that f has odd degree.

Lemma 21.3. *Let $\mathcal{T}$ be a triangulation of B_n, let $\mathcal{K}$ be a triangulation of S_{n-1} that is the restriction of $\mathcal{T}$ to ∂B_n. Let $\mathcal{L}$ be another triangulation of S_{n-1}, and let $\overline{f} : V(\mathcal{T}) \to V(\mathcal{L})$ be a simplicial map. Suppose $f : V(\mathcal{K}) \to V(\mathcal{L})$ is the restriction of $\overline{f}$ to the boundary. Then f has even degree.*

Proof. Let A_n be the chain of all n-dimensional simplices in $\mathcal{T}$. Then A_{n-1}, as in the lemma above, is so that $A_{n-1} = \partial A_n$. Then $\overline{f}_{\#n}(A_n) = 0$, since $\mathcal{L}$ has no n-dimensional simplices. So

$$C_{n-1} = f_{\#n-1}A_{n-1} = f_{\#n-1}\partial A_n = \overline{f}_{\#n-1}\partial A_n = \partial\overline{f}_{\#n}A_n = 0.$$

□

22 Special triangulations*

For $0 \leq k \leq n-1$, we define the hemispheres H_k^+ and H_k^- (of S_{n-1}) by:

$H_k^+ = \{x \in S_{n-1} : x_{k+1} \geq 0, x_{k+2} = ... = x_n = 0\}.$

$H_k^- = \{x \in S_{n-1} : x_{k+1} \leq 0, x_{k+2} = ... = x_n = 0\}.$

Note that H_{n-1}^+ and H_{n-1}^- are "north" and "south", and that the equator is both given by $H_{n-1}^+ \cap H_{n-1}^-$ and by $H_{n-2}^+ \cup H_{n-2}^-$.

We say a triangulation $\mathcal{K}$ of S_{n-1} "respect hemishperes" if it contains a subcomplex that triangulates H_k^+ and H_k^- for all k. A k-chain C_k is antipodal to a k-chain D_k if

$$D_k = \{-\sigma : \sigma \in C_k\}.$$

If C_k is antipodal to itself, we will call it antipodally symmetric.

Lemma 22.1. *Let $\mathcal{K}$ be an antipodally symmetric triangulation of S_{n-1} that respects hemispheres. Let $\mathcal{L}$ be an antipodally symmetric triangulation of S_{n-1} and suppose $f : V(\mathcal{K}) \to V(\mathcal{L})$ is an antipodal simplicial map. Then* $\deg(f) = 1$.

Proof. For each k, denote by A_k^+ the k-chain of all k-simplices in $\mathcal{K}$ contained in H_k^+, similarly A_k^-. Let $A_k = A_k^+ + A_k^-$ (this is, in fact, a union). Then

$$\partial A_k^+ = \partial A_k^- = A_{k-1}, \text{ for } 1 \leq k \leq n-1.$$

Let $C_k^+ = f_{\#k}(A_k^+)$, $C_k^- = f_{\#k}(A_k^-)$ and $C_k = C_k^+ \cup C_k^-$. Then

- $C_k = f_{\#k}(A_k)$, and
- $\partial C_k^+ = \partial f_{\#k}(A_k^+) = f_{\#k-1}\partial(A_k^+) = f_{\#k-1}(A_{k-1}) = C_{k-1}$, and similarly for ∂C_k^-.

By Lemma 21.2, we just need to show that $C_{n-1} \neq 0$. Suppose it is. Then $C_{n-1} = C_{n-1}^+ + C_{n-1}^-$ implies $C_{n-1}^+ = C_{n-1}^-$. We now proceed by induction:

- C_{n-2} is the boundary of an antipodally symmetric chain.

  We know that A_{n-1}^+ is antipodal to A_{n-1}^-, and f is antipodal, so $C_{n-1}^+ = f_{\#n-1}(A_{n-1}^+)$ is antipodal to $C_{n-1}^- = f_{\#n-1}(A_{n-1}^-)$. So C_{n-1}^+ is antipodally symmetric.

  Assume $k \leq n-2$, and $C_k = \partial D_{k+1}$ for some antipodally symmetric $(k+1)$-chain D_{k+1}.

- C_{k-1} is the boundary of an antipodally symmetric chain.

  Let $D_{k+1} = E_{k+1} + E'_{k+1}$, where E'_{k+1} is antipodal to E_{k+1}. Then

$$C_k^+ + \partial E_{k+1} = C_k^- + \partial E'_{k+1}.$$

  Since C_k^+ is antipodal to C_k^-, and E_{k+1} is antipodal to E'_{k+1}, it follows that $D_k = C_k^+ + \partial E_{k+1}$ is antipodally symmetric. Also,

$$\partial D_k = \partial C_k^+ + \partial\partial E_{k+1} = C_{k-1}.$$

Therefore C_0 is the boundary of an antipodally symmetric 1-chain, but $C_0 = f_{\#0}A_0^+ + f_{\#0}A_0^-$ is a pair of antipodal points. The boundary of an antipodally symmetric 1-chain consists of an even number of pairs of antipodal points (end points of intervals). Hence C_0 cannot be the boundary of an antipodally symmetric 1-chain. A contradiction, which ends the proof. □

Proof of Tucker's lemma. By the previous lemma λ must have odd degree, but by Lemma 21.3, λ has even degree. Contradiction. □

Finally, we now show how Tucker's Lemma and Borsuk-Ulam's Theorem are equivalent.

Tucker $\implies$ BU: Suppose there exists $f : B_n \to S_{n-1}$ which is antipodal on the boundary. Let $\delta > 0$ be so that whenever x and x' in B_n are at distance $\leq \delta$, we have

$$\|f(x) - f(x')\|_\infty \leq \frac{2}{\sqrt{n}}.$$

(such a δ exists because f is continuous, and defined on a compact set, so Heine-Cantor Theorem asserts that f is uniformly continuous.)

Note that $\|y\|_\infty \geq 1/\sqrt{n}$ for each $y \in S_{n-1}$.

Now, choose a triangulation $\mathcal{T}$ of B_n that respects hemispheres, is antipodal on the boundary, and has simplex diameter $< \delta$ (meaning, $x, x' \in \sigma \in \mathcal{T}$ implies $d(x, x') < \delta$.) Then, define $\lambda : V(\mathcal{T}) \to \{\pm 1, ..., \pm n\}$ by first making

$$\kappa(v) = \min \left\{ i : |f(v)_i| \geq \frac{1}{\sqrt{n}} \right\},$$

and then

$$\lambda(v) = \begin{cases} +\kappa(v) & \text{if } f(v)_{\kappa(v)} > 0, \\ -\kappa(v) & \text{otherwise.} \end{cases}$$

Then, as f is antipodal on the boundary, it follows that λ is antipodal in S_{n-1}. Thus, by Tucker's Lemma, there exists $\sigma = \{v, v'\} \in \mathcal{T}$ with $\lambda(v) = i$ and $\lambda(v') = -i$, but then $d(v, v') < \delta$ while

$$f(v)_i \geq \frac{1}{\sqrt{n}} \text{ and } f(v')_i \leq \frac{-1}{\sqrt{n}}.$$

Thus $\|f(v) - f(v')\|_\infty \geq \frac{2}{\sqrt{n}}$, contradicting the choice of δ.

- BU $\implies$ Tucker: Assume λ is a simplicial map from a triangulation $\mathcal{T}$ of B_n to $\diamond_{n-1}$ that is antipodal on the boundary. Then the affine extension of λ is a continuous map from B_n to S_{n-1}, which is antipodal on the boundary.

23 Kneser graphs

A *Kneser graph* with parameter (n, r), denoted by $K_{n:r}$, is the graph whose vertices are the r-subsets of a set on n elements, with two r-subsets adjacent if they are disjoint. As we have already seen, $K_{n:1}$ is precisely the complete graph K_n, and $K_{n:r}$ is the empty graph on $\binom{n}{r}$ vertices provided $r > n/2$. If $r = n/2$, then $K_{n:r}$ is a matching on $\binom{n}{n/2}$ vertices. When $r = 2$, it is a useful mental exercise to convince yourself that $K_{n:2}$ is the complement of the line graph of K_n (in particular, $K_{5:2}$ is the well-known Petersen graph, that you should research in case you haven't met it before).

In 1955, Kneser conjectured that the chromatic number of $K_{n:r}$, for $n \geq 2r$, is $n-2r+2$. It is relatively easy to find a colouring with this many colours: let $S \subseteq [n]$ be a vertex in $K_{n:r}$; if this vertex contains an element of the set $[n-2r]$, colour S with the minimum element it contains; if not, then the elements of S are taken from a set of size precisely $2r$, thus the induced subgraph of all such vertices is a matching, which can be colourable with two colours.

Exercise 23.1. Draw the Petersen graph, and colour it according to the above scheme.

Our goal is to show that the upper bound $\chi(K_{n:r}) \leq n - 2r + 2$ is actually tight. This result was first shown by Lovász in 1978 using topological methods, and further simplified by Bárány in the same year, making use of the Borsuk Ulam Theorem and a lemma due to Gale. Before, we try to motivate the importance of Kneser graphs.

Let G be a graph, and let M be a matrix whose rows are indexed by vertices, and columns indexed by all independent sets of G. The chromatic number of G is the solution to the following optimization problem

$$\begin{aligned} \min \quad & \mathbf{1}^{\mathsf{T}} y \\ \text{subject to} \quad & My \geq \mathbf{1} \\ & \mathbf{0} \leq y \leq \mathbf{1} \\ & y \text{ integral.} \end{aligned}$$

That is, y is a 01 vector which indicates which columns of M (independent sets) will be chosen to cover all vertices (the all 1s vector on the right hand side of the inequality), and the objective function being minimized makes it in such way that you will pick the least possible number 1s in y, that is, the least possible number of independent sets to cover $V(G)$.

Upon dropping the integrality condition, the problem becomes less restricted, thus its optimum value decreases. This new optimum is called the *fractional chromatic number*, and is denoted by $\chi_f(G)$. Thus, $\chi_f(G) \leq \chi(G)$.

The above optimization program (with the integrality constraint removed) admits a dual program, given by

$$\begin{aligned} \max \quad & \mathbf{1}^{\mathsf{T}} x \\ \text{subject to} \quad & M^T x \leq \mathbf{1} \\ & \mathbf{0} \leq x \leq \mathbf{1}. \end{aligned}$$

Note that if integrality constraints are added to this program, the vector x will denote a choice of vertices no two of which belonging to the same independent set, thus x will determine a clique in the graph, in fact, the maximization will give the that the optimum of the above program with integrality constraints is $\omega(G)$ (size of the largest clique in G). Without the integrality constraints, we have its fractional version, given by $\omega_f(G)$. Note trivially that $\omega(G) \leq \omega_f(G)$. More importantly, both programs above have been called duals from one another because of the following inequality.

Exercise 23.2. Prove that

$$\omega_f(G) \leq \chi_f(G).$$

(In fact, equality holds, but you wouldn't be able to prove that so easily. You are however invited to try, using a similar technique to what we used when we discussed von Neumann's minimax theorem.)

Thus, we have the chain of inequalities $\omega \leq \omega_f \leq \chi_f \leq \chi$. Moreover, with $\alpha(G)$ denoting the size of the largest independent set in G, another important inequality is obtained by noticing that $x = \mathbf{1}/\alpha(G)$ is a feasible solution to the maximization problem given above, hence

$$\frac{n}{\alpha} \leq \omega_f \leq \chi.$$

Exercise 23.3. Explain, combinatorially, why $\omega \leq \chi$ and $n/\alpha \leq \chi$.

Now, what all that have to do with Kneser graphs? Kneser graphs provide examples of graphs which can make $\chi_f \leq \chi$ (and thus both combinatorial inequalities above) ridiculously weak.

To see that, we will produce a feasible solution y to the minimization formulation for $\chi_f(K_{n:r})$. Consider the independent sets of $K_{n:r}$ which consist of those r-subsets which contain a given element. There are n of such independent sets, each of size $\binom{n-1}{r-1}$. Let y be the vector (whose entries are indexed by the independent sets of $K_{n:r}$) which assign value $1/r$ to each of these special independent sets, and 0 to the remaining. Because each r-subset, vertex of $K_{n:r}$, is contained in precisely r of these special independent sets (one for each of its elements), it follows that y is feasible to the minimization formulation of χ_f, and its value is n/r. Hence $\chi_f(K_{n:r}) \leq n/r$.

Exercise 23.4. Research the famous Erdős-Ko-Rado Theorem, and show that equality actually holds.

By choosing $n = 3r - 1$, we have that $\chi_f(K_{n:r}) \leq 3$ (in this case, it is also very easy to see that $\omega(K_{n:r}) = 2$.) However, Lovász-Kneser Theorem shows that $\chi(K_{n:r}) = r + 1$.

Another important connection between Kneser graphs and fractional chromatic numbers comes from the somewhat surprising result that, for any given graph G,

$$\chi_f(G) = \min\{n/r : G \to K_{n:r}\},$$

where $G \to K_{n:r}$ indicates that there exists a graph homomorphism from G to $K_{n:r}$. (A graph homomorphism is a function that maps vertices from one graph to vertices of the other in such a way that any edge is mapped to an edge.) You are invited to attempt showing the above result (or research it), showing first that if $G \to H$, then $\chi_f(G) \leq \chi_f(H)$. If that is too hard, try at least convincing yourself that

$$\chi(G) = \min\{n : G \to K_n\}.$$

(This one is easy.)

24 Chromatic Number of Kneser Graphs

Our goal is to show that $K_{n:r}$ cannot be coloured with fewer than $n-2r+2$ colours.

We follow the path due to Bárány. This proof uses the last alternative formulation of the Borsuk-Ulam Theorem, due to Lyusternik– Shnirel'man, which we restate below.

Theorem 24.1 (Lyusternik– Shnirel'man). *For any cover $U_1, ..., U_{n+1}$ of the sphere S_n by $n+1$ open sets, there is at least one set containing a pair of antipodal points.*

The idea is to distribute the n elements used to define the r-subsets that form the vertices of $K_{n:r}$ nicely on the sphere S_d with $d = n-2r$. Then, use a colouring of the r-subsets with fewer than $n-2r+2$ colours to produce a covering with open sets of the sphere. Then and finally, with the theorem above, show that the colouring used could not have been a proper colouring of $K_{n:r}$.

The key step in the above scheme consists in being able to distribute the n points nicely on the sphere. In order to achieve that, the following lemma will be used.

Lemma 24.2 (Gale 1956). *For every $d \geq 0$ and every $k \geq 1$, there exists a set $X \subseteq S_d$ of $2k+d$ points such that every open hemisphere of S_d contains at least k points of X.*

Proof. Let $n = 2k+d$. Let $a_1, ..., a_n$ be distinct real numbers, and consider the following matrix, defined with a choice of an integer m that satisfy $n \geq 2m+2$.

$$G = \begin{pmatrix} 1 & 1 & \dots & 1 \\ a_1 & a_2 & \dots & a_n \\ \vdots & \vdots & \vdots & \vdots \\ a_1^{2m} & a_2^{2m} & \dots & a_n^{2m} \end{pmatrix}.$$

Our first claim is that the rows of G are linearly independent. This is easy to see, as the existence of a nonzero vector v so that $v^\mathsf{T} G = 0$ implies that each a_i is a root of a polynomial of degree at most $2m$ whose coefficients are the entries of v. Clearly this cannot be the case.

Now let $x \neq 0$ be so that $Gx = 0$. We want to show that x has a balanced number of negative and positive entries. Assume x has at most m negative entries. Consider

$$g(t) = \prod_{x_i < 0} (t - a_i).$$

This has degree at most m, and thus $f(t) = g(t)^2$ has degree at most $2m$. Thus, there exists a vector y so that

$$y^\mathsf{T} G = (f(a_1), ..., f(a_n)).$$

Note that $y^\mathsf{T} G$ is ≥ 0, and that its ith entry is 0 if and only if $x_i < 0$. Since $y^\mathsf{T} G x = 0$, it follows that x can have no positive entries, and thus $Gx = 0$ implies the existence of m linearly dependent columns in G, a contradiction to the fact that its rank is $2m+1$. The conclusion is that x has at least $m+1$ negative entries, and upon considering $-x$ instead, we also conclude that x has at least $m+1$ positive entries.

Choose $m = k-1$. Note that $n - (2m+1) = d+1$. Let N be the $n \times (d+1)$ matrix whose columns are a basis for the kernel of G. The rows of N form a set of n vectors in

$\mathbb{R}^{d+1}$, so that for any vector $a \in \mathbb{R}^{d+1}$, we have that Na has at least k positive entries (after all, Na belongs to the kernel of G.) Thus, for any a, the open half-space determined by a contains at least k rows of N. Upon normalizing the rows of N, we have found the points in S_d sought by the theorem. □

Theorem 24.3. *let $n \geq 2r$. The chromatic number of $K_{n:r}$ is at least $n - 2r + 2$.*

Proof. Distribute n points on the sphere S_d, with $d = n-2r$, so that every open hemisphere of S_d contains at least r of these points. Consider a colouring of the vertices of $K_{n:r}$ with $d+1 = n-2r+1$ colours, and define open sets $A_1, ..., A_{d+1}$ of S_d by making $x \in A_i$ if there is at least one r-set in the open hemisphere defined by x which has been coloured with colour i. Note that these open sets cover the sphere.

By Theorem 24.1, there exists an i and $x \in S_d$ so that $x \in A_i$ and $-x \in A_i$, but this implies that there is an r-set of the points in $H(x)$ and another in $H(-x)$. These are disjoint, thus adjacent in $K_{n:r}$, but must have forcefully received the same colour. □

In 2002, Joshua Greene (an undergrad!) found a simplification in the above proof that entirely bypasses the need of Gale's Theorem, with the added benefit of providing a strengthening on Theorem 24.1. This is the topic of the following exercise.

Exercise 24.4.

(a) Prove that whenever S_d is covered by $d+1$ sets $A_1, ..., A_{d+1}$, each either open or closed, then there is an i and $x \in S_d$ so that $x \in A_i$ and $-x \in A_i$ (this is for those of you that have studied topology in $\mathbb{R}^n$.)

(b) Set $d = n - 2r + 1$ and place n points in S_d so that no subset of $d + 1$ of these points belong to the same hyperplane through the origin. Explain why this is possible.

(c) Consider a colouring with d colours. Define sets A_i, with i from 1 to d, by having $x \in A_i$ if and only if the open hemisphere defined by x contains an r-set that has been coloured with colour i. Let A_{d+1} be the complement of the union of the A_is. Apply (a) to this collection of sets. Conclude that if the x belonging to A_i and $-A_i$ is so that $i \leq d$, we are done.

(d) Assume now $x \in A_{d+1}$ and $-x \in A_{d+1}$. Conclude that there is a hyperplane containing at least $d + 1$, a contradiction.

A second important remark is that Matoušek found a proof of the Lovász-Kneser theorem that applies Tucker's lemma directly to obtain the lower bound on the chromatic number, thus providing a purely combinatorial proof of said theorem.

Finally, Schrijver provided a remarkable strengthening of Lovász-Kneser. Consider as before r-subsets of $[n]$. Such a subset will be called stable if it does not contain any two adjacent elements modulo n. In other words, S corresponds to an independent set in the cycle graph C_n, whose vertices have been labelled 1 to n. The Schrijver graph $SG_{n:r}$ is the induced subgraph of $K_{n:r}$ whose vertices correspond to the stable sets in $[n]$.

Theorem 24.5. *The chromatic number of the Schrijver graph with parameters n and r, for $n \geq 2r$, is $n - 2r + 2$.*

To prove this theorem, we leave the following strengthening of Gale's theorem as an exercise.

Exercise 24.6. Prove that there exists a set of $2r+d$ points in S_d, naturally identified with $[2r+d]$, so that every open hemisphere contains at least one stable r-subset. (Hint: get creative on the proof of Gale's theorem.)

Index