Sperner's Lemma and its applications

Shagnik Das and Tibor Szabó

Extremal Combinatorics, FU Berlin, WiSe 2017–18

1 Introduction

We now begin our study of the topological method in combinatorics, which dates back to 1978, when Lovász used the Borsuk–Ulam Theorem to prove Kneser's Conjecture. This landmark proof, which turns 40 this November, inspired many other uses of topology to prove combinatorial results, making topology an indispensible component of the modern combinator's toolkit.

That being said, we should not take this for granted, and it is worth taking a moment to reflect just how remarkable it is that topology can be used in combinatorics at all.[1] After all, Wikipedia defines topology as being "concerned with the properties of space that are preserved under continuous deformations."[2] How, then, does one apply it to study discrete objects?

There are two ways to proceed here. One can take a combinatorial problem, embed it in a continuous setting, and then apply a topological result, as Lovász did in his aforementioned proof. Alternatively, one can prove a combinatorial version of a topological theorem, and then apply it directly to the discrete problems at hand. It is this second approach we shall take in this chapter, where we study Sperner's Lemma. Although all the statements herein may at first sight appear completely discrete, a more careful look under the surface will reveal their true topological nature.[3]

2 Sperner's Lemma for flat-earthers

Vaguely speaking, Sperner's Lemma states that if a triangulated d-dimensional simplex is coloured nicely, then we can always find a mutlicoloured simplex. In order for this to be a meaningful mathematical statement, we need to define these terms precisely, and we shall do so gradually, restricting ourselves to the familiar two-dimensional setting in this section.

2.1 What does the lemma say?

Within the comfort of two dimensions, Sperner's Lemma concerns triangulated triangles, which we now describe. We start with a triangle $T = v_1 v_2 v_3$ with vertices v_1, v_2 and v_3. We obtain a triangulation $\mathcal{T}$ by dividing T into smaller triangles. Formally, we add vertices and edges in such a fashion that each internal face has exactly three vertices on its boundaries. Moreover, any two of the smaller triangles must either be disjoint, intersect in a vertex, or intersect in a *common* edge. Below we see a couple of illustrative[4] examples.

[1] Although this is a slippery slope — for instance, it is incredible that you are a collection of tens of trillions of cells working together in near-perfect harmony to ensure that you can read these notes, typed by another set of tens of trillions of cells and delivered to you via a network of billions of computers, that is mostly used by companies with unfathomable amounts of money to harvest your personal information and send you advertisements tailored to your interests. Pondering the miraculous nature of our very existence can quickly descend into a catatonic state of wonder.

[2] Here, too, one should be a little careful — Wikipedia's definition of combinatorics is "an area of mathematics primarily concerned with counting ... and certain properties of finite structures." However, the topology quote is perfect for my intended purposes.

[3] One might say that Sperner's Lemma is discreetly topological.

[4] Or so I hope.

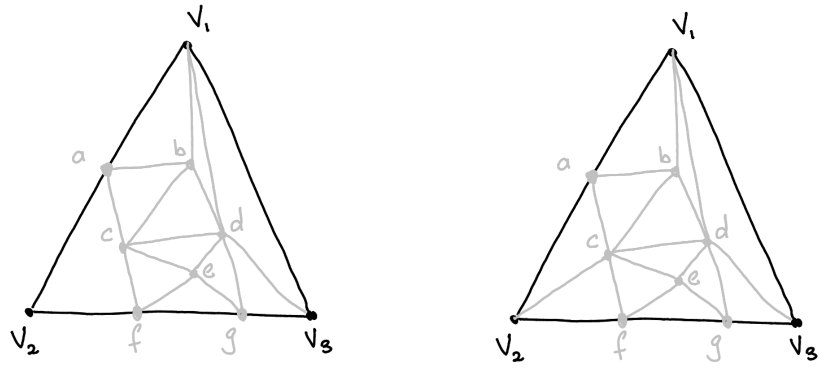

Figure 1: Two divisions of the triangle spanned by v_1, v_2 and v_3.

In the two divisions above, the triangle $T = v_1v_2v_3$ has been subdivided into smaller triangles. However, the division on the left is not a valid triangulation, because the bottom-left face has four vertices on its boundary — v_2, a, c and f. This issue also manifests itself in the fact that the intersection between the triangles v_2af and abc, namely the edge ac, is an edge of abc but not an edge of v_2af. This issue is resolved in the division on the right by further dividing v_2af into two further triangles, v_2ac and v_2cf. This division of T is indeed a triangulation.

In Sperner's Lemma, we "colour" the vertices of a triangulation $\mathcal{T}$ with the "colours" 1, 2 and 3. The goal is to find a multicoloured triangle in the subdivision, which is a triangular face whose vertices all receive different colours. This is clearly not always possible — for instance, if we colour each vertex 1, we cannot hope to find a multicoloured triangle. Indeed, the lemma only applies to the so-called *Sperner colourings*, where we impose restrictions on the colours of vertices on the boundary of T:

- For $1 \leq i \leq 3$, the vertex v_i is coloured i.
- For $1 \leq i < j \leq 3$, any vertex on the side of T between v_i and v_j is coloured either i or j.

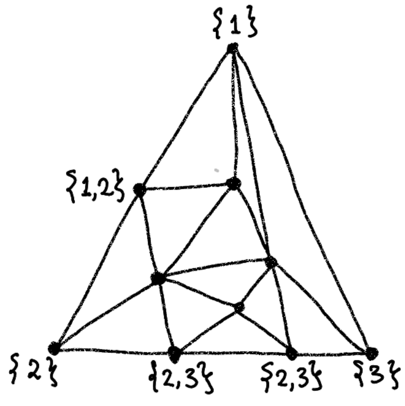

Figure 2: A triangulation, and the colours permitted on the boundary vertices.

Provided these conditions are satisfied, Sperner's Lemma asserts that we will always find a small multicoloured triangle.

Theorem 2.1 (Sperner's Lemma in two dimensions)**.** *Let $T = v_1v_2v_3$ be a triangle, and let $\mathcal{T}$ be a triangulation of T. There is then a multicoloured triangle in every Sperner colouring of $\mathcal{T}$.*

Before we don our serious-mathematician-hats and discuss a proof of this theorem, let us observe that this is quite a remarkable statement, as we are placing no restrictions whatsoever on the colours of the interior vertices. Indeed, had you not encountered this statement within a theorem environment in these course notes, but instead imagined that it was whispered to you by a marmoset clutching a piece of bread, then you might well have thought it impossible that such a statement could be true.

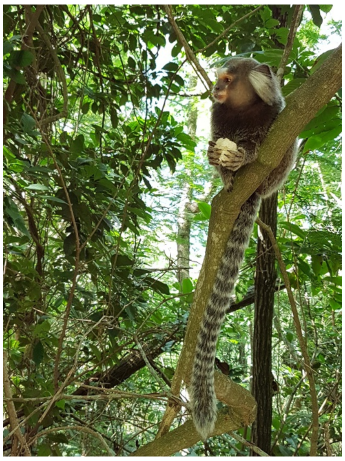

Figure 3: A marmoset, probably contemplating how Sperner's Lemma could possibly be true.

In this case, faced with such doubts, you would surely spend some time trying to find a counterexample.[5] Let us now save you that time by providing below a couple of Sperner colourings of a triangulation, in which we certainly do find multicoloured triangles (shaded).

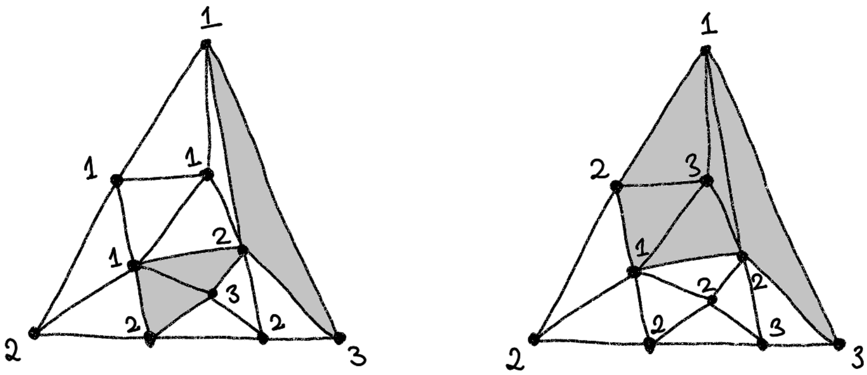

Figure 4: Some evidence in favour of Sperner's Lemma.

Naturally, this "evidence" is not intended to win over your mathematical brains — after all, there are infinitely many possible triangulations, with most admitting many different Sperner colourings, so what can we possibly learn from two small examples? — but rather to sway your mathematical hearts.[6] Now that you are beginning to believe Sperner's Lemma might possibly be true, we shall proceed with a proof to show that it is.

[5]To be fair, any mathematical musings with monkeys are likely to be the result of a fever dream, and you could be forgiven for first tending to your temperature before colouring lots of triangulations.

[6]A cynic might claim the only purpose of this page was to shoehorn a picture of a marmoset into these notes. We shall not dignify these accusations with a response.

2.2 A two-dimensional proof

The proof of Theorem 2.1 is the kind of proof students dream of — short,[7] elegant, and memorable.

Proof of Theorem 2.1. Given the triangulation $\mathcal{T}$ of the triangle $T = v_1v_2v_3$, fix a Sperner colouring $\varphi : V(\mathcal{T}) \to \{1, 2, 3\}$ of the vertices of the triangulation. We now build a graph G in the following fashion:

- We create a vertex $x_F \in V(G)$ for each triangular face F in $\mathcal{T}$. We add one extra vertex x_T representing the external face (that is, everything outside T).
- Two vertices x_F and $x_{F'}$ are adjacent in G if and only if F and F' share an edge, one of whose vertices is coloured 1 and the other coloured 2.

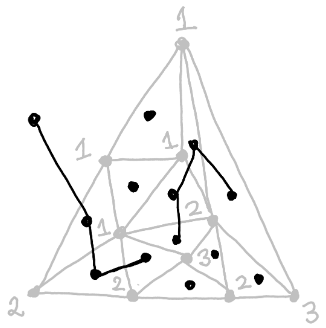

Figure 5: A Sperner colouring of a triangulation (in grey) and the associated graph (in black).

Let us examine the degrees of the vertices in G. By construction, the degree of a vertex x_F (or x_T) is the number of edges on the boundary of F (or on the boundary of T) whose endpoints are coloured 1 and 2.

We start with the degree of x_T. What can we say about the number of edges[8] on the boundary of T whose endpoints are coloured 1 and 2? The second property of the Sperner colourings implies that the only place we can find such edges in $\mathcal{T}$ is between v_1 and v_2. Let $u_1, u_2, \dots, u_t$ be the vertices appearing on this side of T, starting from $u_1 = v_1$ and ending at $u_t = v_2$.

Now consider the sequence of colours $(\varphi(u_i))_{i=1}^t$ of these vertices. Again, by the second property of Sperner colourings, we must have $\varphi(u_i) \in \{1, 2\}$ for each i. Every time we have $\varphi(u_{i+1}) \neq \varphi(u_i)$, we have an edge $\{u_i, u_{i+1}\}$ on the boundary of T whose vertices get the colours 1 and 2, corresponding to an edge containing x_T in G. By the first property of a Sperner colouring, $\varphi(u_1) = \varphi(v_1) = 1$ and $\varphi(u_t) = \varphi(v_2) = 2$. Thus the sequence must change an odd number of times, meaning that the degree of x_T is odd.

From the Handshake Lemma, we know that there must be an even number of vertices of odd degree in a graph G. Since x_T has odd degree, it follows that there must be another vertex, which must be of the form x_F for some internal face $F = uvw$, of odd degree.

Since x_F has positive degree in G, it must have an edge with the colours 1 and 2. Without loss of generality, suppose $\varphi(u) = 1$ and $\varphi(v) = 2$. If $\varphi(w) \in \{1, 2\}$, then there would be precisely one other edge of F receiving both colours 1 and 2, and so x_F would have degree two in G. Hence we must have $\varphi(w) = 3$, and so F must be a multicoloured triangle, as required. $\square$

We remark that this proof actually shows something stronger — any Sperner colouring of a triangulation of a triangle contains an odd[9] number of multicoloured triangles. Indeed, the

[7] So much so that one could question whether the result can really be of any use, and yet here we are, waxing lyrical for pages on end.

[8] Recall that in making the triangulation $\mathcal{T}$, we could add arbitrarily many vertices on the boundary of T, so this number of edges can be unbounded.

[9] And, in particular, positive.

Handshake Lemma shows that there are an odd number of internal vertices of odd degree, and in the last paragraph we proved that any such vertex corresponds to a multicoloured triangle.

2.3 The game of Hex

We round up this section by showing that even this two-dimensional version of Sperner's Lemma has some remarkable applications. Our example of choice is the game of Hex, which has nothing to do with witchcraft, but everything to do with the hexagonal lattice.[10] We first give a quick overview of the rules.[11]

The game is played on an $n \times n$ hexagonal grid, with two players, Joker and Nyssa,[12] alternately claiming empty hexagonal cells by inscribing their initials in them, continuing until all cells in the grid have been filled. Joker wins if he has claimed a path of cells from the left end of the board to the right, while Nyssa wins if she has a path from the top of the grid to the bottom.[13] An example with $n = 5$ is shown below, with Nyssa's winning path highlighted.

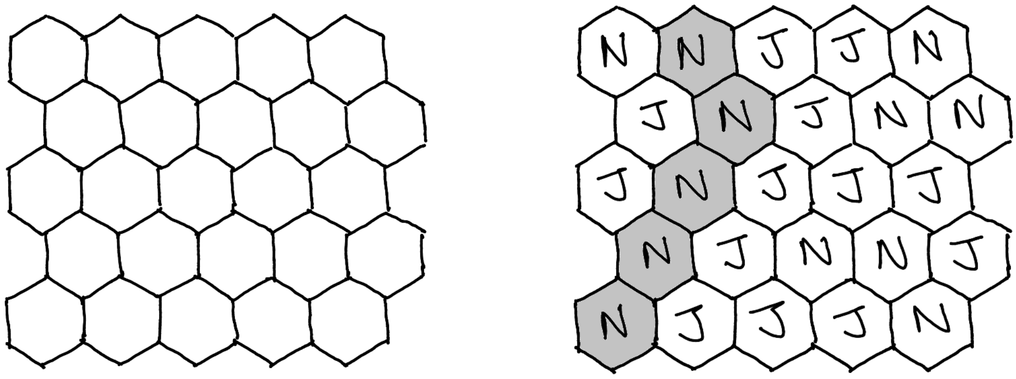

Figure 6: A 5×5 hexagonal lattice, and the board at the end of a game.

Before we go any further, we would be remiss if we did not comment on the history of this game. Although first invented by Piet Hein in 1942, it was independently introduced by John Nash in 1948.[14] John Nash was a remarkable academic, winning both the Nobel Prize for Economics for his fundamental work in game theory[15] and the Abel Prize for his research on nonlinear partial differential equations.[16] That kind of résumé is one that even Hollywood could not ignore, with Russell Crowe depicting the eponymous Nash in *A Beautiful Mind*.[17]

[10] When it comes to board games, humans tend to prefer a square grid, but I imagine bees lounging in their honeycomb would play on hexagons.

[11] In humanity's golden age, this game would have required no introduction, but with the advent of reality TV and Snapchat, nobler pastimes such as Hex have regrettably fallen by the wayside.

[12] I had originally planned on naming the players John and Nash, for reasons that will be apparent in one paragraph's time. However, I then realised that this might cause confusion between John and Nash, the players, and John Nash, the mathematician after whom they were named. As my illustrations were already drawn, I was locked into my choice of initials, and was preparing this during a dark night, so ... .

[13] Usually one would end the game as soon as a winning path is completed, but the analysis is more straightforward if we assume all the cells have been claimed. This does not affect the outcome of the game — a fun exercise (with many solutions) is to show that it is impossible for both players to create winning paths.

[14] The name "Hex" comes from the version of game marketed by the Parker Brothers in 1952. Hein called the game either "Con-tac-tix" or "Polygon," while Nash's version was apparently called "John" by his classmates at Princeton, since it could be played on the hexagonal tiles found in bathrooms. With such clever wordplay, it is no wonder that these bright young minds were snapped up by Princeton!

[15] Those familiar with the subject will no doubt know of Nash already, for he lends his name to Nash equilibria.

[16] At time of press, he is the only person to have received both of these prizes.

[17] This film's success may well have a spawned a series of films about mathematicians — we have recently been treated to *The Imitation Game* (Turing), *The Man Who Knew Infinity* (Ramanujan) and *The Theory of Everything* (Hawking).

As is perhaps befitting one of the fathers of Game Theory, Nash appears to have had quite the affinity for board games.[18] Indeed, in a much-quoted scene from the movie,[19] Nash is depicted playing Go against a fellow Princeton student.[20] At the end of the game, he says (spoilers ahead)

> "You should not have won. I had the first move, my play was perfect ... the game is flawed."

We shall now prove that Hex is a flawless game — with perfect play, Joker will always win. Some standard game-theoretic arguments apply, which we briefly sketch here.[21] Since this is a finite game of perfect information, either one of the players can guarantee a win, or both players can force a draw. Moreover, by strategy stealing, Nyssa (the second player) cannot have a winning strategy. Indeed, suppose for contradiction that she does. In this case, Joker can make an arbitrary first move (because an extra cell can never hurt), and then pretend to be second player, using Nyssa's strategy. If Nyssa also uses her winning strategy, they should both win, but this is not possible. Hence, either Joker has a winning strategy, or both players can force a draw. The following corollary of Sperner's Lemma shows that every game of Hex has a winner, and so it follows that Joker should always win. Do note, though, that we have only proved the *existence* of a winning strategy for Joker; we have not constructed one, and indeed, an explicit winning strategy is unknown (except for small values of n).[22]

Corollary 2.2 (The Hex Theorem)**.** *Every game of Hex has a winner.*

Proof. We start by building a graph Γ_n out of the $n \times n$ hexagonal grid. Replace each hexagon H with a vertex v_H, and put an edge between v_H and v'_H precisely when H and H' intersect in an edge. We then add some external vertices: v_1 on the left, v_2 on the top, v_3 on the right, and v_4 on the bottom. Each of the external vertices is adjacent to the vertices of all the hexagonal faces on the corresponding edge of the grid.

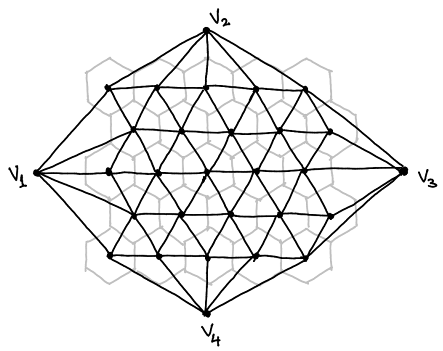

Figure 7: The graph Γ_5.

Now suppose for contradiction that a game of Hex is played without Joker creating a left-right path or Nyssa creating a top-bottom path. We then *label* (not *colour*) each vertex of Γ_n with either a 'J' or an 'N' as follows. The vertices v_1 and v_3 are labelled J, while v_2 and v_4 are labelled N. Every internal vertex v_H is labelled the same way as the corresponding hexagon H in the grid — Joker's cells are labelled J and Nyssa's cells are labelled N.

[18] Aside from Hex, he also created So Long Sucker (warning: profanity).

[19] `https://www.youtube.com/watch?v=GmlSSSN7C78`

[20] Hollywood may not be a bastion of factual accuracy, but as this is just to provide some context for our next result, I am happy to take historical liberties.

[21] For a more thorough handling of this topic, see next year's Algorithmic Combinatorics course.

[22] This explains how Nyssa could win our example above, and why the game is still the most fun you can have on a sunny afternoon.

Given this labelling, we now define a colouring $\varphi : V(\Gamma_n) \to \{1, 2, 3\}$:

$$\varphi(v) = \begin{cases} 1 & \text{if } v \text{ has label J and there is a path from } v \text{ to } v_1\text{, on which all labels are J} \\ 2 & \text{if } v \text{ has label N and there is a path from } v \text{ to } v_2\text{, on which all labels are N .} \\ 3 & \text{otherwise} \end{cases}$$

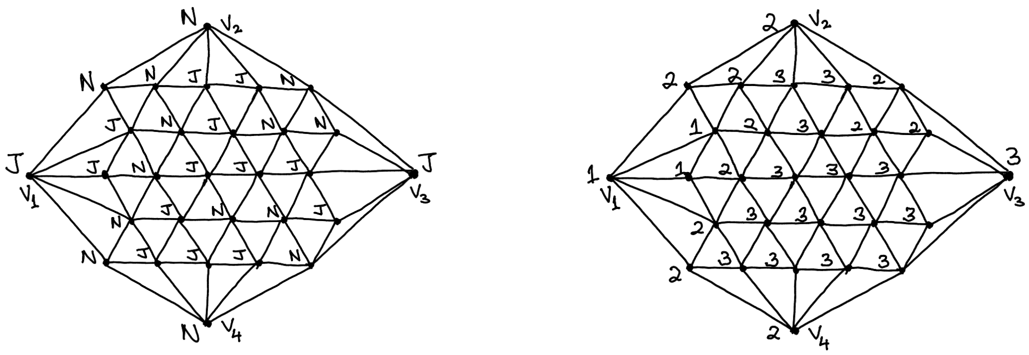

Figure 8: The labelling and colouring of Γ_5 corresponding to a game from Figure 2.3.

We now wish to apply Theorem 2.1 to Γ_n and φ. However, Sperner's Lemma applies to triangulations of triangles, while Γ_n is a triangulation of the square $S = v_1v_2v_3v_4$. We can rectify the situation by "twisting" Γ_n, moving v_4 upwards until it is in line with v_1 and v_3. In this fashion, we can view Γ_n as a triangulation of $T = v_1v_2v_3$.

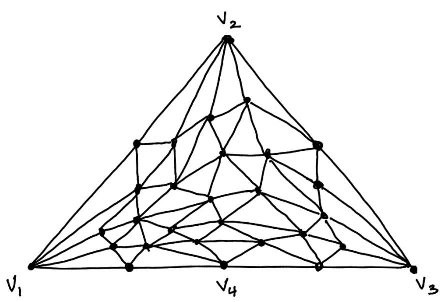

Figure 9: Obtaining a triangulation of a triangle from Γ_5.

Of course, topologically speaking, there is no difference between a square and a triangle.[23] However, we need to be careful about our choice of vertices here, as we need to ensure φ is a Sperner colouring.

For the first property, we trivially have $\varphi(v_1) = 1$ and $\varphi(v_2) = 2$. Since the label of v_3 is J, its colour can only be 1 or 3. If $\varphi(v_3) = 1$, then Joker must have claimed a left-right path, contradicting our assumption that there is no winner. Hence, we must have $\varphi(v_3) = 3$. By a similar argument, $\varphi(v_4) = 3$.

As for the second property, we consider the three sides of T in turn. Let $u_{1,2}$ be the vertex between v_1 and v_2. Since it is adjacent to both v_1 and v_2, it cannot receive the colour 3, regardless of how it is labelled. Hence $\varphi(u_{1,2}) \in \{1, 2\}$. If we let $u_{2,3}$ be the vertex between v_2 and v_3, observe that it corresponds to a hexagon on the right edge of the grid. If it had colour 1, that would imply Joker had built a winning path, contradicting our assumption. Hence $\varphi(u_{2,3}) \in \{2, 3\}$. Finally, we turn our attention to the side between v_1 and v_3. We have already shown $\varphi(v_4) = 3$. The other

[23] People often make fun of topologists, saying that a topologist is someone who drinks coffee out of a doughnut. While that may seem like a harmless joke, discrimination against topologists begins from a very young age.

two vertices on this side correspond to hexagons on the bottom edge of the grid, and since Nyssa has not won the game, they cannot have the colour 2. Thus $\varphi(u) \in \{1, 3\}$ for these two vertices as well.

Thus φ is truly a Sperner colouring of Γ_n, and by Theorem 2.1, there must be a multicoloured triangle, but we will finish the proof by showing this is impossible. Let the vertices of this triangle be u_1, u_2 and u_3, and, without loss of generality, let us assume $\varphi(u_i) = i$ for $1 \leq i \leq 3$. Since $\varphi(u_1) = 1$, u_1 must have the label J, and have a J-labelled path back to v_1. Similarly, u_2 has label N, and an N-labelled path back to v_2. If u_3 had label J, then we could append it to the path of u_1, obtaining a J-labelled path from u_3 to v_1. However, we would then have $\varphi(u_3) = 1$. Thus u_3 must have label N, but then we can append it to the path of u_2, which would imply $\varphi(u_3) = 2$. Hence there cannot be a multicoloured triangle, proving the impossibility of a draw. $\square$

3 Sperner's Lemma in higher dimensions

It is now time to put on your dD glasses, because we are about to see Sperner's Lemma in its full d-dimensional beauty. Fortunately, with all the work we have already invested in the two-dimensional setting, our main task will be to understand what we mean by triangles and triangulations in higher dimensions.

3.1 Definitions in HD

The generalisation of the triangle is, naturally, the d-simplex. The faces of a simplex will also be of importance, and so we define these terms here.

Definition 3.1 (d-simplex). *A* d-simplex T *is the convex hull of* $d+1$ *points* $v_1, v_2, \ldots, v_{d+1} \in \mathbb{R}^d$ *that are affinely independent.*[24] *The points* v_i *are the* vertices *of* T*, and we write* $V(T) = \{v_1, v_2, \ldots, v_{d+1}\}$.

Given a subset $I \subseteq V(T)$ *of size* $k+1$*, the* k*-simplex spanned by* $\{v_i : i \in I\}$ *is said to be a* (k-)face *of* T*, denoted* F_I*. A* facet *is a* $(d-1)$*-face.*

We also required the notion of triangulations, which consisted of decomposing a triangle into smaller triangles. The corresponding object in higher dimensions is a simplicial subdivision, which is a well-behaved decomposition of a simplex into smaller simplices.

Definition 3.2 (Simplicial subdivision). *A* simplicial subdivision *of a* d*-simplex* T *is a finite set of* d*-simplices* $\mathcal{T} = \{S_1, S_2, \ldots, S_m\}$ *such that* $T = \cup_{i=1}^m S_i$*, and, for every* $1 \leq i \neq j \leq m$*,* $S_i \cap S_j$ *is either empty or a common face of both* S_i *and* S_j*. The vertices of a subdivision are the vertices of all its* d*-simplices; that is,* $V(\mathcal{T}) = \cup_{i=1}^m V(S_i)$.

Finally, just as in two-dimensions, we will be interested in colourings of the vertices of our subdivisions, wiith restrictions placed on the vertices of the boundary.

Definition 3.3 (Sperner colourings). *Given a simplicial subdivision* $\mathcal{T}$ *of a* d*-simplex* T*, where* $V(T) = \{v_1, v_2, \ldots, v_{d+1}\}$*, a* Sperner colouring *is a map* $\varphi : V(\mathcal{T}) \to [d+1]$*, such that whenever a vertex* $u \in V(\mathcal{T})$ *lies on the face* F_I *of* T *for some subset* $I \subseteq [d+1]$*, we have* $\varphi(u) \in I$.

A simplex $S_i \in \mathcal{T}$ *is a* multicoloured simplex *if all* $d+1$ *colours appear on the vertices of* S_i.

Note that this implies, in particular, that the vertices of the original simplex T all receive distinct colours: since $v_i \in F_{\{i\}}$, we must have $\varphi(v_i) = i$. Furthermore, any vertex on the edge between v_i and v_j is coloured either i or j, any vertex on the triangle spanned by v_i, v_j and v_k is coloured either i, j or k, and so on.

[24] The points are said to be affinely independent if the vectors $\{v_i - v_{d+1} : 1 \leq i \leq d\}$ are linearly independent in $\mathbb{R}$.

3.2 The lemma in full

Wtih these definitions in place, we can now state Sperner's Lemma[25] in full.

Theorem 3.4 (Sperner's Lemma). *Every Sperner colouring of a simplicial subdivision of a d-simplex contains a multicoloured simplex.*

Observe that Theorem 2.1 is precisely Theorem 3.4 in the case $d = 2$. Some reflection on our previous proof reveals that the very same argument, coupled with induction on the dimension, will suffice to prove the general case as well.[26]

Proof. Using induction on d, we will in fact prove a little more: for all $d \geq 1$, every Sperner colouring of a simplicial subdivision of a d-simplex contains an odd number of multicoloured simplices (and, in particular, at least one).

For the base case, where $d = 1$, observe that a 1-simplex is simply a line segment with endpoints v_1 and v_2, and a simplicial subdivision consists of a sequence of vertices $(u_1, u_2, \ldots, u_t)$, where $u_1 = v_1$ and $u_t = v_2$. We must have $\varphi(u_i) \in \{1, 2\}$ for each i, with $\varphi(u_1) = 1$ and $\varphi(u_t) = 2$. Moreover, a pair (u_i, u_{i+1}) with $\varphi(u_i) \neq \varphi(u_2)$ corresponds to a multicoloured simplex. Since the final colour is different from the first colour, the sequence of colours must switch an odd number of times, giving us an odd number of multicoloured simplices.

For the induction step, let $d \geq 2$, and assume the theorem is true for $(d-1)$-simplices. Let T be a d-simplex with vertices $V(T) = \{v_1, v_2, \ldots, v_{d+1}\}$, $\mathcal{T} = \{S_1, \ldots, S_m\}$ a simplicial subdivision of T, and $\varphi : V(\mathcal{T}) \to [d+1]$ a Sperner colouring of $\mathcal{T}$.

We again build a graph G. As vertices we take $V(G) = \{x_0, x_1, \ldots, x_m\}$, where x_0 corresponds to the $(d-1)$-face $F_{[d]}$ of T, and for $1 \leq i \leq m$, x_i represents the simplex S_i in the subdivision. For $0 \leq i < j \leq m$, the vertices x_i and x_j are adjacent in G if the corresponding simplices intersect in a $(d-1)$-simplex whose vertices are coloured with all colours in $[d]$.

Now observe that the simplicial subdivision $\mathcal{T}$ of T restricts to a $(d-1)$-dimensional simplicial subdivision $\mathcal{T}'$ when restricted to $F_{[d]}$. Moreover, it is easy to see that $\varphi' = \varphi|_{V(\mathcal{T}')}$ is a Sperner colouring of this lower-dimensional subdivision, only using the colours in $[d]$. By the induction hypothesis, $\mathcal{T}'$ contains an odd number of multicoloured simplices, each of which corresponds to an edge in G containing x_0. Hence x_0 has odd degree.

By the handshake lemma, the number of the remaining vertices in G with odd degree must be odd. Now consider a vertex x_i, $1 \leq i \leq m$, of positive degree. The corresponding simplex S_i must therefore have a facet with the colours $\{1, 2, \ldots, d\}$. Say the vertices on this facet are $\{u_1, u_2, \ldots, u_d\}$, where $\varphi(u_j) = j$ for $1 \leq j \leq d$. Let u be the unique vertex of S_i not appearing on this facet. If $\varphi(u) = r \in \{1, 2, \ldots, d\}$, then there is precisely one other facet of S_i with the colours $\{1, 2, \ldots, d\}$; namely, the one spanned by $\{u_1, u_2, \ldots, u_{r-1}, u, u_{r+1}, \ldots, u_d\}$. Hence, in this case, S_i would have even degree.

Therefore, if x_i has odd degree, we must have $\varphi(u) = d+1$, and so S_i is a multicoloured simplex. We therefore have an odd number of multicoloured simplices in $\mathcal{T}$, completing the proof. □

3.3 Where is the topology?

Now that we have seen Sperner's Lemma and its proof, you may be left wondering where the topology, which we made such a fuss about in the introduction, actually is. After all, the proof is purely combinatorial[27] — we build a graph, and look at the degrees of this graph, using nothing more than the Handshake Lemma.

Of course, the topology lies within the notions of simplicial subdivisions and Sperner colourings. The restrictions placed on the colouring depend on which faces the vertices lie, and this is

[25]Not to be confused with Sperner's Theorem, which we encountered earlier. One way to recall the difference is that Sperner's **L**emma is about co**L**ouring simp**L**ices, while Sperner's **T**heorem is about an**T**ichains of se**T**s.

[26]Should we so desire, we could use Theorem 2.1 as the base case. However, to give a self-contained proof, we will start with $d = 1$, and reprove the two-dimensional result.

[27]Of course, one could argue over where the boundary between topology and combinatorics lies. Some people have said that combinatorics is nothing more than the slums of topology. We do not like some people.

information is, in some sense, topological. Furthermore, in applications of Theorem 3.4, such as those we will see later, one often needs a suitable simplicial subdivision, which can be obtained through topological means.

However, the inherent topological nature of Theorem 3.4 is perhaps reflected in the fact that it is equivalent to the well-known Brouwer Fixed Point Theorem.[28]

Theorem 3.5 (Brouwer Fixed Point Theorem). *Let B^d be the d-dimensional ball. Any continuous map $f : B^d \to B^d$ has a fixed point.*

Apart from being a beautiful result, this theorem has some significant real-world consequences. For instance, the case $d = 3$ implies that after a topologist has stirred her morning doughnut of coffee, there is some atom that returns to its initial location.

4 Independent transversals

We shall now discuss a few applications of Theorem 3.4, so as to assure you that the generalisation to higher dimensions was worthwhile. The first of these concerns independent transversals and, to get you even more excited for the topic, we will open with some motivation from the real world.

4.1 Forming committees

There is a university[29] where every decision — the hiring of a new professor, the setting of degree requirements, the choice of font for official communication — must be made by a committee, tasked with ensuring that the decision made is fair, unbiased, and in the best interests of the university.

We require two things of these committees: that they be representative, containing one member from each department, and that they be able to make decisions quickly. Unfortunately, these two goals may well be in conflict with each other. For instance, if everyone in the mathematics department were to hate everyone in the physics department,[30] then any representative commitee would contain a pair of enemies, and it would be difficult to get them to agree on anytihng.

Fortunately, though, the department members are too busy with their work to hate many other faculty members. If nobody hates many other people, and the departments are large, we might hope to be able to form a committee whose members all get along with each other. We now seek to determine whether this is actually true, and to see what ‘many’ and ‘large’ need to mean.

4.2 A graphic reformulation

To make the problem more precise, we shall, of course, restate it in graph-theoretic terms. We build a graph G, with one vertex for each member of the university faculty. The vertices of this graph are partitioned into the different departments, so we have $V(G) = V_1 \dot\cup V_2 \dot\cup \ldots \dot\cup V_r$, with each V_i being the faculty members of a particular department.[31] A representative committee is a choice of one vertex from each part V_i, $1 \le i \le r$, otherwise known as a transversal.

In order to keep track of which pairs of dons hate each other, and should not be selected together to serve in a committee, we put an edge for every pair of vertices representing incompatible people.

[28] Showing the equivalence of Theorem 3.4 and Theorem 3.5 is a rewarding exercise. Bonus fun: the Hex Theorem (Corollary 2.2) is also equivalent to both of these results.

[29] Which, for legal reasons, shall not be named in these notes.

[30] Which is not unthinkable: the mathematicians may well be jealous of all the cool toys the physicists get to play with in their labs, while physicists could resent mathematicians for not even being able to solve simple differential equations describing how viscous fluids flow.

[31] Naturally, a committee had to be formed to decide how the indices would be assigned to the departments. However, as this was a crucial step in solving the problem of making efficient committees, this decision took a long time to make. Eventually, following a bitter struggle that saw no fewer than four professors resigning in disgust, the department of underwater basket weaving won the right to be V_1. The department of mathematics was stuck with V_r, primarily because their representative was preoccupied with doodling graphs on pieces of scrap paper throughout the meetings, and never spoke a word.

We thus want to ensure that the transversal we choose does not span any edges; in other words, it should be an independent set.

Definition 4.1 (Independent transversal). *Let G be an r-partite graph with parts $V_1, V_2, \ldots, V_r$. An* independent transversal *is a set $S \subseteq V(G)$ such that*

(i) S is an independent set, and

(ii) $|S \cap V_i| = 1$ for all $1 \leq i \leq r$.

Our goal is to find conditions on the graph G that imply the existence of an independent transversal. The first thing we might think to require is that the parts be large, as this gives us more freedom to avoid edges when selecting a vertex from each part. For an extreme example, if we had two parts of size one, and those two vertices were adjacent, it would be impossible to build an independent transversal.

However, even large part sizes might not be enough if we have vertices of high degree as well. Sticking with the extreme examples, if our graph G was complete, we would certainly not find an independent transversal. We shall therefore also bound the maximum degree of the graph. Putting these requirements together gives rise to the extremal problem we will be solving.

Question 4.2. *Given natural numbers Δ and r, what is the minimum $p = p(\Delta, r)$ such that every r-partite graph with parts of size at least p and maximum degree at most Δ contains an independent transversal?*

At first sight it is not clear that $p(\Delta, r)$ should exist at all — while making the parts larger provides more freedom of choice, it also provides more opportunities to place edges that destroy independent transversals. However, a natural greedy algorithm shows that $p(\Delta, r)$ is indeed finite.

Claim 4.3. *For all Δ and r, we have $p(\Delta, r) \leq (r-1)\Delta + 1$.*

Proof. We build an independent transversal greedily: starting from V_1 and working up to V_r, we choose a vertex $v_i \in V_i$ that is not adjacent to the vertices $\{v_j : 1 \leq j < i\}$ chosen previously. At every step, we have already chosen at most $r-1$ vertices, each of which has at most Δ neighbours in the current part. Thus, as long as we have at least $(r-1)\Delta + 1$ vertices to choose from, there is always a valid choice, allowing us to complete an independent transversal. $\square$

We now have an upper bound, though it remains to be seen how good it is. What kind of lower bound can we give for $p(\Delta, r)$? We earlier noted that if the graph G was complete, we would not be able to form an independent transversal. For a less wasteful obstruction, if we had a complete bipartite graph $K_{p,p}$ between two parts, that would be enough to prevent an independent transversal. Hence, if $p \leq \Delta$, we cannot guarantee the existence of an independent transversal. This leaves us with the bounds $\Delta + 1 \leq p(\Delta, r) \leq (r-1)\Delta + 1$ whenever $r \geq 2$.[32] Which is closer to the truth?

4.3 The chromatic number

Before answering that question, we shall describe an important application[33] of independent transversals, which will provide us with an alternative construction for the lower bound $p(\Delta, r)$.

Recall that the chromatic number $\chi(H)$ of a graph H is the minimum p such that we can find a colouring $c : V(H) \to [p]$ such that whenever $\{h, h'\} \in E(H)$, we have $c(h) \neq c(h')$. The problem of finding a proper colouring of a graph is a fundamental one in combinatorics, and it can be naturally reduced to finding an independent transversal in a graph, showing the importance of our problem.[34]

[32] Trivially, $p(\Delta, 1) = 1$.

[33] That is, something mathematical, not something from the real world.

[34] In terms of complexity, deciding whether a graph is p-colourable for any $p \geq 3$ is $\mathcal{NP}$-complete (i.e. hard), and therefore so is the problem of finding an independent transversal in a graph.

We build a graph G with vertices $V(G) = \{(h,i) : h \in V(H), i \in [p]\}$. The vertex (h,i) in G represents colouring the vertex $h \in H$ with colour $i \in [p]$. Hence, if we make vertex parts $V_h = \{(h,i) : i \in [p]\}$ for each $h \in V(H)$, a transversal corresponds to a colouring of the vertices in H.

However, we need to find a proper colouring. By defining edges of G appropriately, we shall ensure that proper colourings of H correspond to independent transversals in G. To this end, let $E(G) = \{\{(h,i),(h',i)\} : \{h,h'\} \in E(H), i \in [p]\}$; that is, two vertices in G are adjacent if they correspond to the same colour and a pair of adjacent vertices in H. We have $\Delta(G) = \Delta(H)$, since G is nothing more than p disjoint copies of H. Furthermore, it is easy to see that a proper colouring of H is precisely an independent transversal of G.

Hence, we must have $p \geq \chi(H)$ in order for G to have an independent transversal. If we take $H = K_{\Delta+1}$, we have $\chi(H) = \Delta(H) + 1$, and so it follows that $p(\Delta, \Delta+1) \geq \Delta + 1$. Since $p(\Delta, r)$ is monotone non-decreasing in r,[35] we have a different construction to show $p(\Delta, r) \geq \Delta + 1$ for all $r \geq \Delta + 1$.

4.4 A sharp lower bound

Although we have so far only seen the lower bound $p(\Delta, r) \geq \Delta + 1$, some sporadic examples were found that showed, for certain values of Δ, one could have $p(\Delta, r) \geq 2\Delta$. These constructions seemed to be worst possible cases, leading Bollobás, Erdős and Szemerédi to make the following conjecture in 1975.

Conjecture 4.4. *For all Δ and r, $p(\Delta, r) \leq 2\Delta$.*

As we shall see in this section, not only is this conjecture correct, but it was the correct conjecture to make — years later, in 2006, Szabó and Tardos provided a superior construction to the ones we have seen thus far, proving that $p(\Delta, r) \geq 2\Delta$ for all Δ and $r \geq 2\Delta$.[36]

Theorem 4.5. *For all Δ and $r \geq 2\Delta$, $p(\Delta, r) \geq 2\Delta$.*

Before going through the details of the construction, let us quickly sketch the underlying idea. We earlier showed $p(\Delta, r) \geq \Delta + 1$ by considering the complete bipartite graph $K_{\Delta,\Delta}$. Here, in a transversal, one is forced to take one vertex from each part of the graph, thus inducing an edge.

Our goal is to show the existence of r-partite graphs with parts of size $2\Delta - 1$ that do not admit an independent transversal. The idea is to take disjoint union of several copies of $K_{\Delta,\Delta}$. We then have to cleverly divide the vertices into parts in such a way that if one were to try to construct an independent transversal, one would eventually be forced to choose one vertex from each part of a copy of $K_{\Delta,\Delta}$, thus spanning an edge. We now make this idea precise.

Proof. By monotonicity of $p(\Delta, r)$ in r, it is enough to show the theorem for $r = 2\Delta$, as then $p(\Delta, r) \geq p(\Delta, 2\Delta) \geq 2\Delta$ for all $r \geq 2\Delta$.

For $-(\Delta - 1) \leq i \leq \Delta - 1$, let G_i be a copy of $K_{\Delta,\Delta}$ with vertex classes labelled A_i and B_i. We take our graph G to be the disjoint union of the graphs G_i, which clearly has maximum degree Δ. We shall now divide the vertices of G into r parts of size $2\Delta - 1$ in such a way that any independent transversal would be forced to contain one vertex from A_0 and one from B_0, giving an edge and, with it, ending the proof.

For $0 \leq i \leq \Delta - 2$, fix an arbitrary vertex $a_i \in A_i$. We define the following parts:

(i) For $1 \leq i \leq \Delta - 1$, let $V_i = (A_{i-1} \setminus \{a_{i-1}\}) \cup B_i$.

(ii) Let $V_\Delta = A_{\Delta-1} \cup \{a_i : 0 \leq i \leq \Delta - 2\}$.

[35] For $r' > r$, we can always extend an r-partite graph to an r'-partite graph by adding $r' - r$ parts of p vertices, without adding any edges. Any independent transversal of the larger graph contains an independent transversal of the original one.

[36] This is, in fact, a special case of the Szabó–Tardos result, who gave the best-possible lower bound on $p(\Delta, r)$ for smaller values of r as well.

Observe that this gives Δ parts which partition $A_0 \cup (\cup_{i \geq 1} V(G_i))$.

Claim 4.6. *Any independent transversal of the parts $V_1, V_2, \ldots, V_\Delta$ must contain a vertex of A_0.*

Proof. Let T be an independent transversal of $V_1, V_2, \ldots, V_\Delta$, and for $1 \leq i \leq \Delta$, let $v_i \in V_i$ be the vertex of V_i contained in T. We start by considering v_Δ. If $v_\Delta = a_0$, we are done. Hence we may assume $v_\Delta \in A_i$ for some $i \geq 1$.

Now consider $v_i \in V_i \subset A_{i-1} \cup B_i$. Since T is independent, we cannot have $v_i \in A_i$, as then it would be adjacent to v_Δ. Thus $v_i \in A_{i-1}$.

By an analogous argument, we can deduce that $v_{i-1} \in A_{i-2}$, and so on, until we deduce that $v_1 \in A_0$, proving the claim. $\square$

We have thus partitioned half of the vertices of G into Δ parts, and ensured that any independent transversal contains a vertex from A_0. To complete the proof, we partition the remaining vertices symmetrically, negating indices and exchange the roles of the B_j and A_j.

More precisely, for $-(\Delta - 2) \leq j \leq 0$, fix an arbitrary vertex $b_j \in B_j$. We define the parts:

(i) For $-(\Delta - 1) \leq j \leq -1$, let $V_j = (B_{j+1} \setminus \{b_{j+1}\}) \cup A_j$.

(ii) Let $V_{-\Delta} = B_{-(\Delta-1)} \cup \{b_j : -(\Delta - 2) \leq j \leq 0\}$.

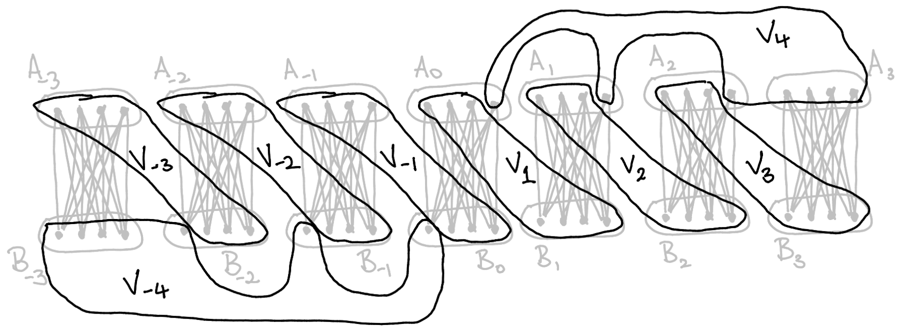

Figure 10: The graph G and its partition when $\Delta = 4$.

The parts $V_{-\Delta}, V_{-(\Delta-1)}, \ldots, V_{-1}$ partition $B_0 \cup (\cup_{j \leq -1} V(G_j))$, and, with essentially the same proof as before, we have the following claim.

Claim 4.7. *Any independent transversal of the parts $V_{-\Delta}, V_{-(\Delta-1)}, \ldots, V_{-1}$ must contain a vertex of B_0.*

To conclude, we view G as a (2Δ)-partite graph with parts $V_{-\Delta}, \ldots, V_{-2}, V_{-1}, V_1, V_2, \ldots, V_\Delta$, each having size $2\Delta - 1$. If this partition were to admit an independent transversal, then by Claims 4.6 and 4.7, this transversal would contain vertices from both A_0 and B_0, which would induce an edge. Hence there is no independent transversal of G, and so $p(\Delta, 2\Delta) \geq 2\Delta$. $\square$

4.5 The upper bound

Theorem 4.5 shows that the upper bound of Conjecture 4.4, if true, would be best possible. As it turns out, Haxell had proven the conjecture a few years earlier in 2001. Together with the previous lower bound, this proves $p(\Delta, r) = 2\Delta$ for all $r \geq 2\Delta$.

Theorem 4.8. *If G is an r-partite graph with maximum degree Δ and parts of size at least 2Δ, then G contains an independent transversal.*

Although Haxell's original proof was algorithmic, she later provided an elegant topological proof using Sperner's Lemma, and it is this second proof that we shall present. The idea is to build a simplicial subdivision of an $(r-1)$-simplex, together with a Sperner colouring, in such a way that a multicoloured simplex corresponds to an independent transversal. To that end, the following definitions will be useful.

Definition 4.9 (Simplicial boundaries)**.** *Given a d-simplex T, its* boundary ∂T *is the union of its facets. Given a simplicial subdivision $\mathcal{T}$ of T, we denote by $\partial\mathcal{T}$ the induced subdivision of ∂T into $(d-1)$-simplices.*

Note that we lose information when we only consider the simplicial subdivision of the boundary — many different subdivisions could have the same boundary subdivision. In particular, there is a standard way to obtain a subdivision from a subdivided boundary: create a new vertex in the interior of T, and add it to each of the $(d-1)$-simplices in $\mathcal{T}'$.

Definition 4.10 (Lifted subdivision)**.** *Give a d-simplex T and a subdivision $\mathcal{T}' = \{S'_1, S'_2, \ldots, S'_m\}$ of the boundary ∂T, the* lifted subdivision $\mathcal{L}(\mathcal{T}') = \{S_1, S_2, \ldots, S_m\}$ *is a subdivision of T, where for some $x_0 \in T \setminus \partial T$, $V(\mathcal{L}(\mathcal{T}')) = V(\mathcal{T}') \cup \{x_0\}$, and $V(S_i) = V(S'_i) \cup \{x_0\}$ for each $1 \le i \le m$.*

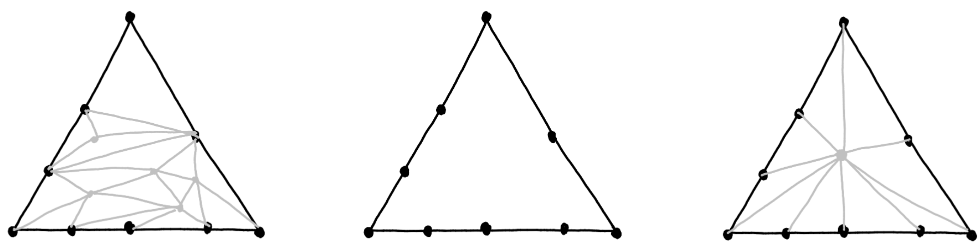

Figure 11: A subdivision $\mathcal{T}$ of T; $\partial\mathcal{T}$ subdivides the boundary ∂T; the lifted subdivision $\mathcal{L}(\partial\mathcal{T})$.

We shall then need to label the vertices of a subdivision $\mathcal{T}$ with vertices from our graph G, with the hope of eventually finding independent transversals. Hence our labelling scheme will have to take into account which parts the vertices are coming from, and which pairs of vertices are adjacent. This gives rise to the following definition.

Definition 4.11 (Independent labelling)**.** *Suppose we have a $(d+1)$-partite graph H with $V(H) = V_1 \dot\cup V_2 \dot\cup \ldots \dot\cup V_{d+1}$. Let T be a d-simplex with vertices $V(T) = \{t_1, t_2, \ldots, t_{d+1}\}$, and let $\mathcal{T}$ b a subdivision of either T or the boundary ∂T. An* independent labelling *is a map $L : V(\mathcal{T}) \to V(H)$ such that:*

(i) for every $I \subset [d+1]$, if a vertex $u \in V(\mathcal{T})$ lies on the face F_I of T, then $L(u) \in \cup_{i \in I} V_i$.

(ii) if a pair of vertices $u_1, u_2 \in V(\mathcal{T})$ are adjacent in the subdivision,[37] *then their labels $L(u_1), L(u_2) \in V(H)$ are* not *adjacent in the graph H.*[38]

The following claim shows that, given a subdivision of the boundary of a simplex, together with an independent labelling of the subdivision, we can extend the subdivision and labelling to cover the whole simplex without changing anything on the boundary.

Claim 4.12. *Let H be a $(d+1)$-partite graph of maximum degree Δ and parts of size at least 2Δ. Let T be a d-simplex, $\mathcal{T}'$ a subdivision of the boundary ∂T, and $L' : V(\mathcal{T}') \to V(H)$ an independent labelling of the vertices in $\mathcal{T}'$. Then there exists a simplicial subdivision $\mathcal{T}$ of T and an independent labelling $L : V(\mathcal{T}) \to V(H)$ such that $\partial\mathcal{T} = \mathcal{T}'$ and $L(u) = L'(u)$ for all vertices $u \in V(\mathcal{T}')$.*

[37]That is, they appear together in a common simplex $S \in \mathcal{T}$.

[38]Note that we allow adjacent vertices in $\mathcal{T}$ to receive the same label, as a vertex in H is not adjacent to itself.

We shall prove this claim in due course, but let us first see how it implies the theorem.

Proof of Theorem 4.8. Let G have parts $V_1 \dot\cup V_2 \dot\cup \dots \dot\cup V_r$, with each part of size at least 2Δ, where Δ is the maximum degree of G. We seek an independent transversal of G.

We will show the existence of a simplicial subdivision $\mathcal{T}$ of an $(r-1)$-simplex T, together with an independent labelling $L : V(\mathcal{T}) \to V(G)$. Given such a labelling, we obtain a colouring $\varphi : V(\mathcal{T}) \to [r]$ by setting, for every $u \in V(\mathcal{T})$, $\varphi(u) = i$ precisely when $L(u) \in V_i$. Property (i) of independent labellings shows that this is, in fact, a Sperner colouring of $\mathcal{T}$.

Hence, by Theorem 3.4, $\mathcal{T}$ contains a multicoloured simplex. By definition of our colouring, that implies that the labels of the vertices of this simplex belong to different parts of G; that is, they form a transversal of G. Moreover, since this vertices lie in a common simplex, they are pairwise adjacent. By Property (ii) of independent labellings, their labels form an independent set. We have thus found the desired independent transversal.

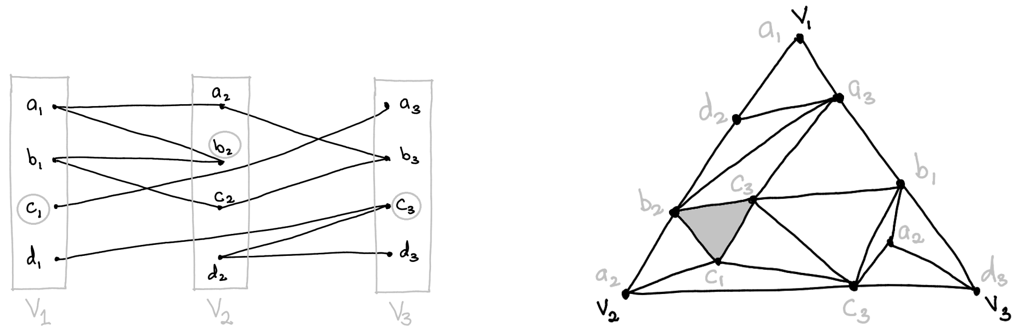

Figure 12: A 3-partite graph and a subdivision of a 2-simplex with an independent labelling. A multicoloured simplex and the corresponding independent transversal are highlighted.

It therefore suffices to find an independently-labelled subdivision of an $(r-1)$-simplex T. Let the vertices of the simplex be $V(T) = \{t_1, t_2, \dots, t_r\}$. We shall build up the desired subdivision and labelling by working our way up through the faces of T, one dimension at a time. We start with the 0-dimensional faces — the vertices. For each vertex t_i of T, choose an arbitrary vertex $v_i \in V_i$ from G, and set $L(t_i) = v_i$.

Now suppose we have independently-labelled subdivisions of all $(d-1)$-dimensional faces of T, for some $1 \le d \le r-1$. For every set $I = \{i_1, i_2, \dots, i_{d+1}\} \subseteq [r]$, consider the d-simplex that is the d-face F_I of T. Let $H = G[\cup_{i \in I} V_i]$ be the $(d+1)$-partite subgraph of G induced by the parts whose indices lie in I.

We then have a subdivision of the boundary of ∂F_I, together with an independent labelling of its vertices with vertices from H. Applying Claim 4.12, we can find a labellling of a d-dimensional subdivision of the face F_I. Since this does not affect the subdivision of the boundary ∂F_I, we can combine these subdivisions of all the d-faces, so that we now have independently-labelled subdivisions of all the d-faces.

When $d = k$, the d-face is the whole simplex T, giving the desired labelled subdivision, and thus completing the proof. ☐

It remains, of course, to prove Claim 4.12. In order to find a nice labelling of a nice subdivision of a simplex, we shall first need the nice subdivision of the simplex. This is given to us by our next claim.[39]

Claim 4.13. *Suppose we have a d-simplex T and a subdivision $\mathcal{T}'$ of the boundary ∂T. There is a subdivision $\mathcal{T}$ of the full simplex T with $\partial \mathcal{T} = \mathcal{T}'$, and an ordering $<$ of $V(\mathcal{T})$, such that:*

[39] It may look like we are proving this à la Zeno — introducing infinitely many claims, each halving the distance to the result — but this is the last claim, promise.

(i) every boundary vertex precedes any interval vertex; that is, if $u \in V(\mathcal{T}')$ and $v \in V(\mathcal{T}) \setminus V(\mathcal{T}')$, then $u < v$, and

(ii) every internal vertex is adjacent in $\mathcal{T}$ to at most $2d$ preceding vertices (internal or boundary).

Given this claim, the proof of Claim 4.12 is straightforward.

Proof of Claim 4.12. We are given $\mathcal{T}'$, a subdivision of the boundary ∂T of the d-simplex T, and an independent labelling $L' : V(\mathcal{T}') \to V(H)$.

By appealing to Claim 4.13, we find a subdivision $\mathcal{T}$ of the simplex T, with $\partial\mathcal{T} = \mathcal{T}'$, together with an ordering $<$ of $V(\mathcal{T})$. Our goal is to create an independent labelling $L : V(\mathcal{T}) \to V(H)$.

For every vertex $u \in V(\mathcal{T}')$, we set $L(u) = L'(u)$. We will process the remaining vertices (those in the interior, $V(\mathcal{T}) \setminus V(\mathcal{T}')$) according to the order $<$.

Suppose we are currently considering some vertex $v \in V(\mathcal{T})$. Since it is in the interior of T, Property (i) in the definition of independent labellings poses no restriction — we could choose a label $L(v)$ from any part of H. However, we shall have to ensure that we maintain Property (ii). In particular, if v is adjacent to some preceding (and thus labelled) vertex u in $\mathcal{T}$, then we cannot have $L(v)$ adjacent to $L(u)$ in H.

However, by Property (ii) of Claim 4.13, there are at most $2d$ such preceding neighbours u of v. Each of these neighbours is adjacent to at most Δ vertices in H, and so there are at most $2d\Delta$ forbidden labels for v. On the other hand, H has $d+1$ parts, each of size at least 2Δ, and so the number of remaining vertices in H is at least $2(d+1)\Delta - 2d\Delta > 0$. Choose one of these arbitrarily to be the label $L(v)$.

Repeating this process for each vertex in turn, we obtain an independent labelling of $\mathcal{T}$. $\square$

4.6 A subdivision of low degeneracy

All we have left, then, in our proof of Theorem 4.8, is to prove Claim 4.13.[40] Now you may wonder why we have roamed into a new subsection of these notes. The truth is that, by reducing Theorem 4.8 to Claim 4.13, we have crossed the border from combinatorics to topology. Unfortunately, we have only been granted short-term tourist visas for these foreign lands, leaving us enough time for a quick sketch without being able to delve into details. Indeed, at some point, we will have to take some topological facts for granted. While this should not live up to your standards of what a proof should be,[41] the aim is to provide you with some intuition for why the claim should be true, and we shall hopefully succeed in this endeavour.

Sketch of the proof of Claim 4.13. The proof is by induction on the dimension, d. The base case, $d = 1$, is straightforward, as we need only add one internal vertex adjacent to the two boundary vertices of the 1-simplex.

Now we proceed with the induction step, with $d \geq 2$. The first observation is that, as previously discussed, we can always lift a subdivision $\mathcal{T}'$ of the boundary ∂T to a subdivision $\mathcal{L}(\mathcal{T}')$ of the simplex T by adding a single internal vertex and adding it to each of the $(d-1)$-simplices in $\mathcal{T}'$. While this is certainly a subdivision, this does not prove Claim 4.13, because this new internal vertex may well be adjacent to too many preceding vertices.

Hence what we do is *imagine* we added all these edges, thus giving a subdivision with a lot of imaginary edges. We then will slowly change this subdivision into something acceptable by removing one imaginary edge at a time. This will require the introduction of many new internal vertices, but, using the induction hypothesis and a little topological wizardry, each of these new internal vertices is adjacent to at most $2d$ previous vertices.

Step 0 Starting with our imaginary subdivision $\mathcal{L}(\mathcal{T}')$ (with only one internal vertex, with an imaginary edge to every boundary vertex), order the imaginary edges arbitrarily.

[40] This is akin to reaching Mount Everest's Base Camp, and remarking that all we have left are 3484 metres.

[41] After all, what follows is not the greatest proof in the world; no, this is just a tribute.

Step 1 In our current subdivision (all real edges + all surviving imaginary edges), consider the next imaginary edge. Call this edge e. In steps 2 to 6, we will modify the subdivision to remove e.

Step 2 In order to remove e, we will have to replace all the simplices containing e in our current subdivision. Consider the union of these simplices — these must account for some volume of space surrounding the edge e. In other words, the union of these simplices looks like[42] a closed d-dimensional ball[43], which we will call B_e, of which e is a diameter.

In three dimensions, you might imagine this union of simplices containing e looks something like a mandarin[44], with each segment of the mandarin representing one of the simplices, and the edge e being the stringy white fibre[45] that runs from top to bottom.

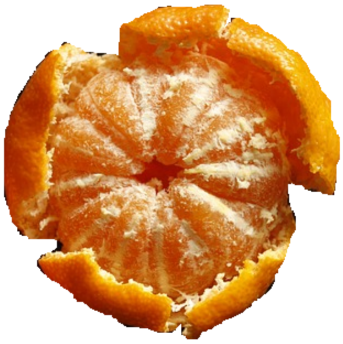

Figure 13: A half-peeled mandarin, or the union of subsimplices containing e?

Step 3 What we need to do is find a new subdivision of B_e that does not use the edge e.[46] We can then use the new subdivision for B_e, and use the rest of the old subdivision for $T \setminus B_e$, and we will have a subdivision of T that does not use e.[47]

Step 4 It is a topological fact that in a nice subdivision,[48] if we consider the union of the simplices that make up B_e, and remove the edge e, what we end up with is essentially the $(d-2)$-dimensional sphere S^{d-2}. In our three-dimensional example, we get an equatorial circle on the surface of the mandarin that separates the endpoints of the edge e.[49]

Step 5 We now consider the $(d-1)$-dimensional closed ball we obtain by "filling in" this sphere S^{d-2}, which we call B'_e. That is, consider the surface of the ball that is exposed when we cut along the equatorial circle. B'_e is topologically equivalent to a $(d-1)$-simplex, and has a subdivided boundary (the remaining faces of the simplices surrounding e).

In Figure 4.6 below, if one looks closely at the first image, one can make out a triangulation of the surface of the mandarin, with each of the segments of the mandarin contributing a triangle. This is not a valid subdivision of B'_e, because the central vertex is in too many edges to boundary

[42] If you smooth out the surface a little, which doesn't really change anything.

[43] The closed unit d-ball B^d is the set of points $\vec{x} \in \mathbb{R}^d$ such that $\sum_{i=1}^d x_i^2 \leq 1$. Its boundary is the $(d-1)$-dimensional sphere S^{d-1}, consisting of all points $\vec{x} \in \mathbb{R}^d$ with $\sum_{i=1}^d x_i^2 = 1$.

[44] Or tangerine, or clementine, or whatever you prefer.

[45] There is presumably a proper term for this, but it shall not be found here.

[46] In our fruity three-dimensional example, we want to find a new way to divide up the interior of the mandarin, rather than splitting it into its segments.

[47] Since our new subdivision for B_e only changes the interior, and not the boundary, we can still combine it with the old subdivision outside B_e.

[48] This is where we drop any pretence of rigour, and shall not attempt to define what we mean by a "nice subdivision." Rest assured that the subdivisions used in this process are all nice enough for this fact to be applicable.

[49] You can think of the endpoints of the edge e as the little green stub of the mandarin and its antipode, situated at the north and south poles with respect to this equatorial circle.

vertices. We fix this by induction: we can add interior vertices to B'_e, with each new vertex being adjacent to at most $2(d-1)$ previous vertices, to obtain a subdivision of B'_e.

Hence the subdivision of the first image represents the initial imaginary subdivision of B'_e in Step 0 of our induction hypothesis. What we will instead use is the final subdivision the induction hypothesis gives us, shown in the second picture.

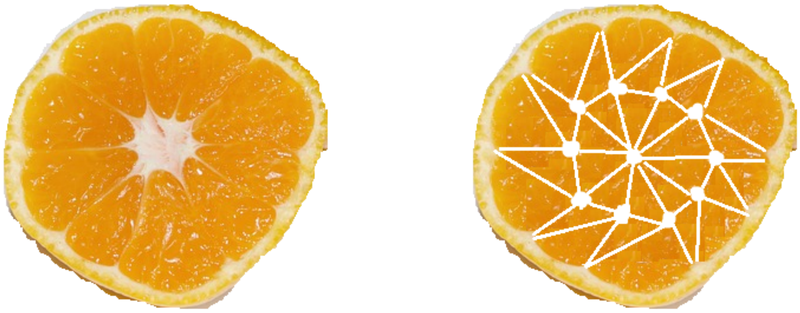

Figure 14: The cut mandarin, exposing the 2-simplex B'_e inside the equator, and the good final subdivision thereof.

Step 6 However, this only gives a subdivision of B'_e, a lower-dimensional simplex inside B_e. What we need is a subdivision of all of B_e.

We do this by lifting each $(d-1)$-simplex in the subdivision of B'_e into two d-simplices — one above the equator and one below. This can be done by adding to each new vertex in B'_e two edges: one to each endpoint of the edge e. This means that instead of being adjacent to at most $2(d-1)$ previous vertices, the new vertices are now adjacent to at most $2d$ previous vertices, but this is still okay. This then provides a subdivision of the whole d-ball B_e into d-simplices, none of which use the edge e.

For our three-dimensional mandarin, this means we find a different way to cut the mandarin. Rather than splitting it into its natural segments, we break it up into different pieces, none of which has a edge that runs the whole length of the mandarin.[50]

Step 7 This new subdivision of B_e, together with the old subdivision of $T \backslash B_e$, gives a subdivision of T that does not use the edge e. Repeating this process for every imaginary edge e, we will obtain a subdivision of T that does not use any imaginary edges.

In this subdivision, every new vertex added was adjacent to at most $2d$ preceding vertices (and the original internal vertex is not adjacent to *any* preceding vertices, since all of those edges have been removed), completing the induction step and the proof of Claim 4.13. □

In case you lost track — we did start many pages ago — this completes our proof of Theorem 4.8, showing $p(\Delta, r) \leq 2\Delta$. If nothing else, you should now appreciate that Sperner's Lemma is truly an example of topological combinatorics!

5 Fair division

For our final application of Sperner's Lemma, we once again stray from our combinatorial cocoon, venturing this time into the fields of economics.[51] One of the fundamental driving forces behind economic theory[52] is the scarcity of resources — that our wants and needs exceed supplies. This

[50] Which is a terrible way to eat a mandarin, in our opinion.

[51] After all, as we have seen, if you want to win a Nobel Prize as a mathematician, economics is your best bet. What are the odds on Satoshi Nakamoto getting the nod?

[52] And, arguably, all human endeavour.

leads to various interesting problems regarding optimal allocation of resources, and we shall focus on that of fair division. There are many examples one could use to illustrate this problem, ranging from the judgment of King Solomon[53] to the division of post-war Berlin.[54] However, rather than relying on historical examples that, while significant, are far removed from your daily lives, we shall instead discuss an important dilemma that people face annually — how should one cut one's birthday cake?

5.1 A piece of cake

Let us set the scene: a young lad named Chris happened to be born an integer number of years ago, and has invited his friends Alfred and Batman over for his birthday party. They do what one usually does at a birthday party — reenact their favourite hat problems, play Coup, and eat Hawaiian pizza. At long last, though, the highlight of the party is brought out from the kitchen. It is time for the birthday cake.

The song is sung, a wish is made, and the candles are blown out. Now, however, comes the tricky part. the cake must somehow be cut and divided between the three friends. Unfortunately, we humans have a tendency to compare what we have to what others have, and then want what they have instead.[55] As we are at a birthday party, though, the goal is to have all the guests be happy, and to prevent any fights from breaking out.[56]

“Sounds simple enough," you think, “just cut the cake into three equal pieces."

That assumes, though, that there is some measure on the cake, so that we can determine when two piecees are *equal*. It further assumes that each friend has the same measure, so that they can all agree that the pieces made are of the same worth. However, neither of these assumptions are ones we wish to make.

We will instead work with a much more general model. Assuming there are n people between whom the cake is to be shared, all we will assume about the people is that, after the cake has been cut into n pieces, each person can identify their favourite piece (or pieces, in case of a tie). We shall *not* assume that the people can rank all n pieces in order of preference, or that they can somehow quantify how much more they prefer their favourite piece to any other piece.[57] We can now more explicitly state what our goal is.

Definition 5.1 (Envy-free division). *When cutting a cake into n pieces, and distributing these pieces between n people, the division is said to be* envy-free *if every person feels that their own piece is as least as good as that of anyone else.*

Of course, it is not always possible to find an envy-free division. Suppose, for instance, that somewhere in your cake is an indivisible yet deadly iota of poison. Assuming further that all of your friends desire to stay alive,[58] then in any division, the person who is served the poisonous piece will no doubt be envious of everyone else's slice. We therefore will need to impose some requirements on our friends' cake preferences, but, as we shall see, these requirements are very mild.

First, though, we need to introduce some notation. We will restrict ourselves to a very particular way of cutting the cake - by making $n-1$ parallel cuts, we shall make n slices. To measure the slices, we use the standard cake units: the cake itself is one cake long, and we shall say that the ith slice is x_i cakes wide. Hence any such division of the cake maps to the point $(x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$,

[53] `http://en.wikipedia.org/wiki/Judgment_of_Solomon`

[54] As students in Berlin, you should be somewhat familiar with this example, but just in case: `http://en.wikipedia.org/wiki/Berlin_Wall`.

[55] As they say, the grass is always greener on the other side, which tends to turn one green with envy.

[56] After all, if you are cutting a cake, you have a cake knife nearby, and so to prevent unwanted bloodshed, it is important to keep ill-will at a minimum. It might, therefore, have been a better idea to have only played Coup *after* the cake was cut.

[57] We are really being very accommodating hosts here, asking the bare minimum we need of our guests in order to please them.

[58] So that they may celebrate their own future birthdays with you.

where $x_i \geq 0$ for all i and $\sum_i x_i = 1$. The collection of all such points forms an $(n-1)$-dimensional simplex.

Definition 5.2 (Cake simplex)**.** *Given $n \geq 2$, the* cake simplex *is the $(n-1)$-dimensional simplex $\Delta_{n-1} = \{(x_1, x_2, \ldots, x_n) : x_i \geq 0, \sum_i x_i = 1\} \subseteq \mathbb{R}^n$.*

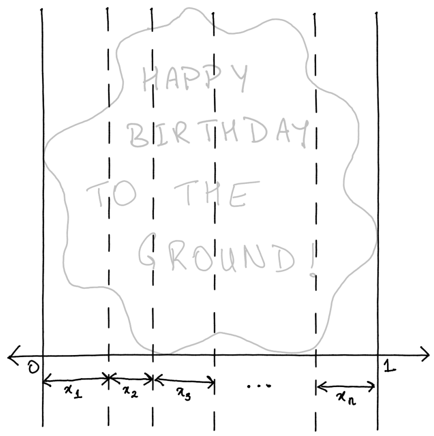

Figure 15: A point in the cake simplex represents a division of the cake.

We can now define two properties that we shall require our friends' preferences to satisfy.

Definition 5.3 (Positive and continuous preferences)**.** *Given a division of the cake represented by $\vec{x} = (x_1, x_2, \ldots, x_n)$, a person's* preference *is the choice of which slice (or slices, if some are equally desirable) that person would choose. We say that their preference is* positive *if they never choose a slice with $x_i = 0$,*[59] *and that their preference is* continuous *if, given some sequence of divisions $\vec{x}^{(1)}, \vec{x}^{(2)}, \ldots \to \vec{x}^*$, if the person prefers the ith slice in each $\vec{x}^{(k)}$, then they also prefer the ith slice in $\vec{x}^*$.*[60]

The following theorem of Simmons and Su shows that, under these two very reasonable assumptions, an envy-free division always exists.

Theorem 5.4. *Suppose we are dividing a cake between n friends. If the preferences of the friends are positive and continuous, then there is an envy-free division $\vec{x}^* \in \Delta_{n-1}$.*

We should note that there are many other results on cake-cutting in the literature, which we sadly do not have time to review. However, the above theorem has some of the mildest restrictions, and is therefore one of the most applicable.

5.2 The Simmons–Su protocol

The proof of Theorem 5.4 is, in a theoretical sense, algorithmic, and this process is known as the Simmons-Su protocol. It relies heavily on Sperner's Lemma, applied to carefully-chosen subdivisions of the cake simplex Δ_{n-1}. These are the so-called barycentric subdivisions.

Definition 5.5 (Barycentric subdivisions)**.** *The* barycentric subdivision *of a one-dimensional simplex is obtained by dividing the 1-simplex into two smaller 1-simplices of equal length, by adding the midpoint of the simplex as a new vertex.*

[59] Informally, the person likes cake, and would prefer any part of the cake to an empty slice.

[60] Informally, this is saying that infinitesimally small changes in the division should not affect one's preferences. Note that in the limit, there might be other slices which are just as preferable, but the person should always be happy with the ith slice of $\vec{x}^*$.

The barycentric subdivision $\mathcal{T}$ *of a d-simplex T is obtained by first obtaining a subdivision $\mathcal{T}'$ of the boundary ∂T by taking the barycentric subdivision of each of the facets. We then lift this to obtain $\mathcal{T} = \mathcal{L}(\mathcal{T}')$, where the internal vertex is placed at the barycentre of T.*[61]

Once we have a subdivision $\mathcal{T}$ of the simplex, we can of course obtain a finer subdivision, by subdividing each simplex in $\mathcal{T}$.

Definition 5.6 (kth barycentric subdivision)**.** *Denote by $\mathcal{T}_1$ the previously-defined barycentric subdivision of a d-simplex T. For $k \geq 2$, the* kth barycentric subdivision *$\mathcal{T}_k$ of T is obtained by taking the union of the barycentric subdivisions of each d-simplex in the $(k-1)$st barycentric subdivision $\mathcal{T}_{k-1}$ of T.*

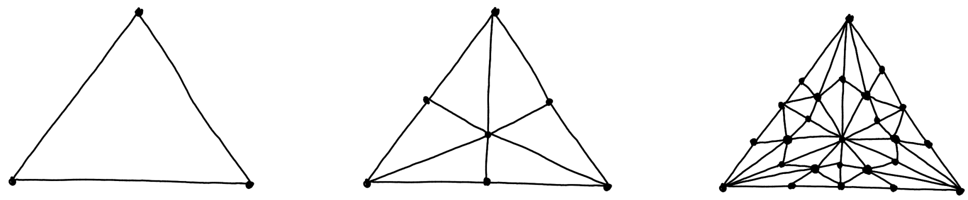

Figure 16: The 2-simplex Δ_2; its barycentric subdivision; its second barycentric subdivision.

We require a couple of properties of the barycentric subdivisions. The first shows that there is a special labelling of the vertices of the kth barycentric subdivision $\mathcal{T}_k$ of Δ_d.

Claim 5.7. *Given the k-th barycentric subdivision $\mathcal{T}_k$ of Δ_d, and a set of $d+1$ symbols $A = \{\alpha_1, \alpha_2, \ldots, \alpha_{d+1}\}$, there is a labelling $\lambda : V(\mathcal{T}_k) \to A$ such that each simplex in $\mathcal{T}_k$ receives all $d+1$ symbols on its vertices.*

Proof. We first consider the case $k = 1$, where we can use the following labelling: given a vertex $u \in V(\mathcal{T}_1)$, we set $\lambda(u) = \alpha_i$ if and only if u lies on an $(i-1)$-face of Δ_d, but not on an $(i-2)$-face. It can be seen (by induction, for instance) that every d-simplex in $\mathcal{T}_1$ contains, for each $0 \leq k \leq d$, exactly one barycentre of a k-face of Δ_d as a vertex, and so this labelling has the desired property.

Now, given $k \geq 2$, recall that the simplices of $\mathcal{T}_k$ are obtained by barycentrically subdividing the simplices in $\mathcal{T}_{k-1}$. By the $k = 1$ case, each of these barycentric subdivisions has a good labelling. Moreover, since the label given to a vertex depends only on the dimension of the face it lies in, vertices that are in several simplices of $\mathcal{T}_{k-1}$ always receive the same label. Hence, given any $u \in V(\mathcal{T}_k)$, fix a simplex $S \in \mathcal{T}_{k-1}$ such that $u \in S$.[62] We then assign the same label to u as it would receive in the labelling of the barycentric subdivision of S given in the case $k = 1$. Since every simplex $S' \in \mathcal{T}_k$ is contained within some simplex of $S \in \mathcal{T}_{k-1}$, and the labelling is good for S, it follows that the vertices of S' receive all different labels. □

The second property we shall need, which we will not prove in these notes, is that the simplices of the barycentric subdivision have smaller diameter than that of the original simplex.

Fact 5.8. *Let T be a d-simplex, and $\mathcal{T}$ the barycentric subdivision of T. Then*

$$\max_{S \in \mathcal{T}} \operatorname{diam}(S) \leq \frac{d}{d+1} \operatorname{diam}(T).$$

With these preliminaries, we can now prove the existence of envy-free divisions.

[61] The *barycentre* is the centre of mass of the simplex. In particular, the barycentre of the cake simplex Δ_{n-1} is the point $(1/n, 1/n, \ldots, 1/n)$. Note that this is more a geometric subdivision than just a topological one — the placement of the vertices, and, as a result, the lengths of the edges play an important role.

[62] Note that we do not require $u \in V(\mathcal{T}_{k-1})$ — there are vertices $u \in V(\mathcal{T}_k) \setminus V(\mathcal{T}_{k-1})$ that we also need to label.

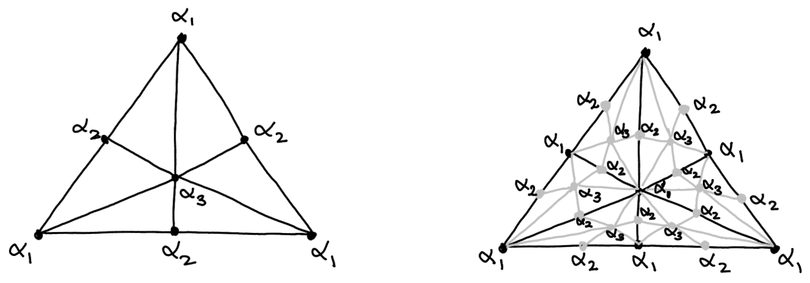

Figure 17: The labellings of $\mathcal{T}_1$ and $\mathcal{T}_2$ of Δ_2.

Proof of Theorem 5.4. Consider the kth barycentric subdivision $\mathcal{T}_k$ of the $(n-1)$-dimensional cake simplex Δ_{n-1}. By Claim 5.7, there is a labelling $\lambda : V(\mathcal{T}_k) \to \{\alpha_1, \ldots, \alpha_n\}$ of the vertices, such that each simplex in $\mathcal{T}_k$ has distinct labels on all its vertices.

We shall now define a Sperner colouring $\varphi : V(\mathcal{T}_k) \to [n]$. Each vertex $u \in V(\mathcal{T}_k)$ receives a label $\lambda(u) = \alpha_i$, say. Moreover, as a point in the cake simplex, $u = \vec{x}$ represents a division of the cake. We shall then ask the ith person which slice of the cake they would prefer, if the cake were divided according to $\vec{x}$, and let $\varphi(u)$ be their answer.

To see that this is a Sperner colouring, observe that for any $I \subseteq [n]$, the face F_I of the simplex Δ_{n-1} is precisely the set of points $\{(x_1, x_2, \ldots, x_n) \in \Delta_{n-1} : x_j = 0 \text{ if } j \notin I\}$. By the assumption of the positivity of preferences, nobody would choose a slice of zero width, and hence the answers for vertices on the face I will always lie within I. Therefore φ is indeed a Sperner colouring.

By Theorem 3.4, it follows that we can find a multicoloured simplex $S \in \mathcal{T}_n$. This means that at each vertex of S, a different answer was given. By the property of our labelling λ, these answers were given by different people. Hence, this multicoloured simplex represents the n different people choosing n different slices of the cake, as we would want in an envy-free division.

However, we are unfortunately not done.[63] While we did get different choices of slices from the n people, these were choices made at different vertices, which correspond to different divisions. This does not give us *one* division where everyone opts for different pieces.

To find such a division, we employ a limiting argument. We have shown above that in each barycentric subdivision $\mathcal{T}_k$, we find a multicoloured simplex $S^{(k)}$. There are n vertices and n colours, giving $n!$ ways in which a simplex can be multicoloured. From this infinite sequence of multicoloured simplices, we pass to an infinite subsequence of multicoloured simplices $(S^{(k_i)})_{i\geq 1}$ that are all coloured the same way. Recall that the colouring encodes the choices of slices given by the people. Without loss of generality, we may assume that the first person always chooses the first slice, the second person always chooses the second slice, and so on.

Let $\vec{x}^{(i)}$ be the vertex of $S^{(k_i)}$ with the label α_1 (that is, it is the vertex for which the first person chose the slice). Since the simplex Δ_{n-1} is compact, this sequence must have a convergent subsequence; suppose we have $\vec{x}^{(i_j)} \to \vec{x}^* \in \Delta_{n-1}$. Since the first person always chooses the first slice for each division $\vec{x}^{(i_j)}$, the continuity assumption implies this person is also content with the first slice at the division $\vec{x}^*$.

Now consider the rth person, for some $2 \leq r \leq n$, and let $\vec{y}^{(i)}$ be the vertex of $S^{(k_i)}$ with label α_r (that is, the vertex for which the rth person chose the slice). Since $x^{(i)}$ and $y^{(i)}$ both lie in $S^{(k_i)}$, we have $\|\vec{x}^{(i)} - \vec{y}^{(i)}\| \leq \operatorname{diam}(S^{(k_i)})$.

By Fact 5.8, it follows that $\operatorname{diam}(S^{(k_i)}) \leq \left(\frac{n-1}{n}\right)^{k_i} \operatorname{diam}(\Delta_n)$, which tends to zero as i tends to infinity. In particular, we must thus have $\vec{y}^{(i_j)} \to \vec{x}^*$ as well. Using the continuity of the rth person's preferences, the rth person is content with the rth slice in the division $\vec{x}^*$.

In conclusion, each of the n people would be happy with a different slice in the division given by $\vec{x}^*$, and hence this is an envy-free division. $\square$

[63]This was evident from the fact that we have not used the continuity assumption on the preferences, nor have we used anything about k.

5.3 Rental harmony

We close these notes with a brief discussion about further applications of the Simmons-Su protocol: sometimes a cake is not just a cake.[64] We have already touched upon some other situations where this could have been applied; for instance, in deciding where to draw the boundaries between the Soviet, British, American and French quarters of Berlin after the Second World War. In more modern examples, the protocol could be used to divide prize money between a winning team, to assign credit for a group project, and so on — the possibilities are endless.

A slightly different setting, which may be of particular interest to you, is the question of rental harmony. Here we have n roommates, moving into an apartment with n different rooms. There are some decisions to be made here — who should move into which room, and what portion of the rent should they pay? We can think of each point $\vec{x} \in \Delta_{n-1}$ as representing a division of the rent between the different rooms, and hope to find an envy-free division — an assignment of rent to the rooms, such that each roommate would prefer a different room.

However, here there is an issue with our assumption of positivity — if there was a room with zero rent assigned to it, it would be quite popular! Hence the preferences, under this mapping of rent divisions to the simplex Δ_{n-1}, would not satisfy the positivity, and thus we could not deduce the existence of an envy-free division.

The solution is to consider not the simplex Δ_{n-1}, but instead its dual. We shall not get into details here, but the interested reader is invited to consult Su's paper.[65] When looking to apply these methods yourself, if you do not have the patience to compute infinitely many subdivisions of the cake simplex, then you'll be happy to hear there's an app for that: Spliddit.

In closing, you may wonder why the Simmons-Su protocol hasn't ushered an era of world peace. If we can divide resources in an envy-free manner, why do governments struggle to pass budgets? Why do we even need governments? Sadly, Simmons and Su did not solve the problem of human greed — while everyone involved may believe their piece is at least as good as everyone else's, that does not stop them from wanting more.

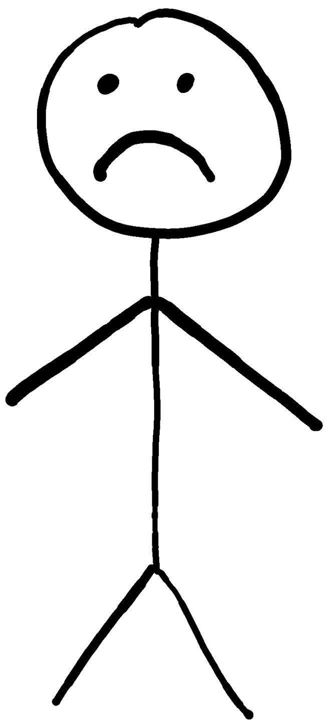

[64] Sometimes, the cake is a lie.
[65] `https://www.math.hmc.edu/~su/papers.dir/rent.pdf`