Topological Methods in Combinatorics

Talk by Pablo Soberón

Notes by Sanjana Das

September 15, 2023

Today we'll discuss combinatorial problems that can be solved by topological methods — especially ones where the topological methods are a bit unexpected.

§1 Colorful Radon

We'll start with the colorful Radon theorem.

Theorem 1.1 (Colorful Radon, Lovász 1992)

Imagine that we are given $d+1$ pairs of points in $\mathbb{R}^d$, where each pair is of a different color. Then there is a colorful partition — a partition of the points into two sets such that each part contains one point of each color — such that the convex hulls of the two parts intersect.

For example, if $d = 2$, we have three pairs of points in the plane; we want to split them into two triangles (each consisting of one point of each color) such that the triangles intersect.

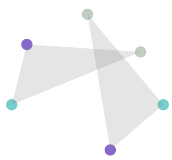

It's possible to solve this using linear algebra, but we'll see the topological proof that Lovász had.

Proof. We'll sometimes use $d = 2$ for illustration (so when we say 'triangle' we really mean 'simplex' in the general case).

We start out with a set of points in $\mathbb{R}^d$, and the first thing we'll do is represent them in $\mathbb{R}^{d+1}$ — we create $d+1$ axes, one for each pair. We have the two blue points correspond to two opposite points on the first axis (which are symmetric about the origin; we remember which blue point in $\mathbb{R}^d$ corresponds to which blue point in $\mathbb{R}^{d+1}$), the two purple points correspond to two opposite points on the second axis, and so on. (This doesn't depend on how our original points were arranged.) Then any colorful triangle in our configuration in $\mathbb{R}^d$ corresponds to some colorful triangle in $\mathbb{R}^{d+1}$.

Now we can look at the set of *all* possible colorful triangles in $\mathbb{R}^d$, and the corresponding set of all colorful triangles in $\mathbb{R}^{d+1}$ — the latter set forms the boundary of an octahedron $\mathbb{O}^d$, which if you squint a bit is essentially the same as a d-dimensional sphere $\mathbb{S}^d$.

So this gives us a continuous function $f\colon \mathbb{S}^d \to \mathbb{R}^d$ (where for each point on our octahedron $\mathbb{O}^d$, we take the corresponding point inside the triangle in $\mathbb{R}^d$, i.e., the same linear combination of our three points in $\mathbb{R}^d$). Now we can use the Borsuk–Ulam theorem.

Theorem 1.2 (Borsuk–Ulam)
If $f\colon \mathbb{S}^d \to \mathbb{R}^d$ is continuous, then there exists $x \in \mathbb{S}^d$ with $f(x) = f(-x)$.

But in this case, if we take two antipodal points x and $-x$ in $\mathbb{S}^d$, they will correspond to complementary triangles, and their f-values are points inside these triangles; so this gives us a colorful partition into intersecting triangles. $\square$

§2 Necklace Splitting Theorem

We'll now discuss the necklace splitting theorem, due to Golberg and West (1986).

Question 2.1. Suppose we're given an open necklace, with pearls strung together on a segment; these pearls have various colors (e.g., blue, yellow, and green), such that there are m colors and there is an even number of pearls of each color.

Now imagine two thieves steal such a necklace, and they want to distribute it such that each thief gets exactly half of the pearls of every color.

What's the minimum number of cuts necessary to do so (i.e., to evenly divide the necklace)?

Of course, the thieves can always divide the necklace by just cutting between all pairs of pearls, but they'd like to be more efficient (i.e., to use fewer cuts).

First, what would a 'bad' necklace look like? One example is if we have all the blue pearls grouped together, then all the green pearls grouped together, then all the yellow pearls, and so on; in this case, we'll need at least m cuts (since we need to cut somewhere inside the stretch of each color).

One 'bad' necklace is if we have all the blue pearls grouped together, all the green pearls grouped, and all the yellow pearls grouped; this means we always need at least m cuts.

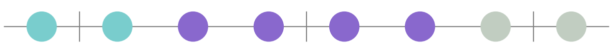

Theorem 2.2 (Necklace splitting theorem)
It is always possible to find a fair partition with at most m cuts.

We'll actually prove a slightly different result, essentially a continuous version of this problem — we'll treat our necklace as the interval $[0, 1]$. Instead of having m colors of pearls, we'll have m absolutely continuous measures on $[0, 1]$.

Theorem 2.3 (Hobby, Rice 1965)

Given m absolutely continuous measures $\mu_1, \ldots, \mu_m$ on $[0, 1]$, we can find a fair partition of $[0, 1]$ (i.e., one where the two pieces have equal measure under each of $\mu_1, \ldots, \mu_m$) using m cuts.

These two statements are actually equivalent — we can approximate a measure by a finite set of points, and we can approximate a finite set of points by measures concentrated near those points.

Now we'll prove this continuous version.

Proof. In the previous problem, we looked at the space of all possible colorful simplices, and parametrized it as a nice space — the surface of an octahedron, which was basically a sphere. We'd like to do something similar here — to parametrize cutting $[0, 1]$ into m pieces. So let's look at the lengths $x_1, \ldots, x_{m+1}$ of these pieces, which sum to 1. If we just looked at these lengths, then the set of such vectors $(x_1, \ldots, x_{m+1})$ would be the m-dimensional simplex Δ^m. But we actually want to keep track of a bit more — we also want to know which thief gets each piece. So we also give these pieces signs — we call our thieves A and B, and we assign a piece $+$ if A gets it and $-$ if B does. So now we have numbers $y_1, \ldots, y_{m+1}$ (where $y_i = \pm x_i$), and the space of all possible partitions can be described as

$$\{(y_1, \ldots, y_{m+1}) \in \mathbb{R}^{m+1} \mid |y_1| + \cdots + |y_{m+1}| = 1\}.$$

This set is again an octahedron $\mathbb{O}^m \cong \mathbb{S}^m$.

Now we have a sphere, so we just need to make a function to use Borsuk–Ulam on. We have m measures, so we can define

$$f(y) = (\mu_1(A), \ldots, \mu_{m+1}(A))$$

to keep track of how much of each measure A gets. Now Borsuk–Ulam gives that there exists a point y with $f(y) = f(-y)$. But $f(-y) = (\mu_1(B), \ldots, \mu_{m+1}(B))$, since flipping all the signs corresponds to flipping who gets each piece. So this point y corresponds to a fair partition. □

§3 Baskets and Fruits

Question 3.1. We're given n baskets, each of which has some amount of each of k kinds of fruit. We want to choose some of these baskets, and we have two goals:

- We want to have at least half of each kind of fruit.
- We want to have as few baskets to carry as possible.

What's the minimum number of baskets we need?

So we have a bunch of baskets with some amount of bananas, blueberries, and so on; and we want to choose as few baskets as possible such that we get at least half of each fruit.

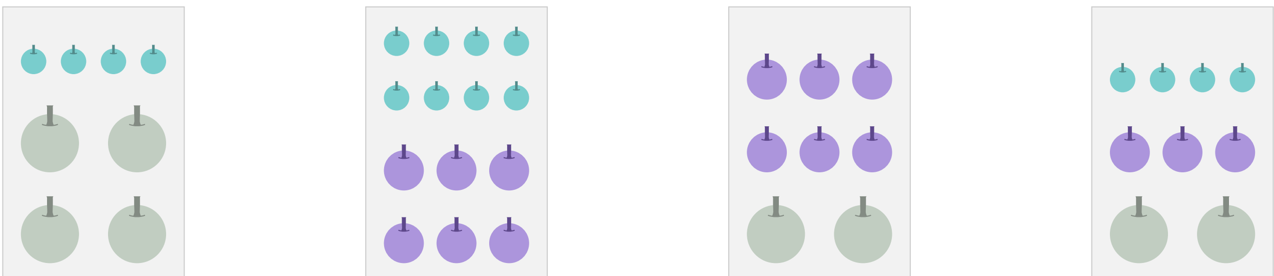

We'll assume that $n \not\equiv k \pmod 2$ — otherwise, when $n \equiv k \pmod 2$, our solution may be off by 1.

First, what would be a 'bad' set of baskets and fruits? We might have the first $k-1$ kind of fruits each filling one basket, and the kth kind evenly distributed among the remaining $n-k$ baskets.

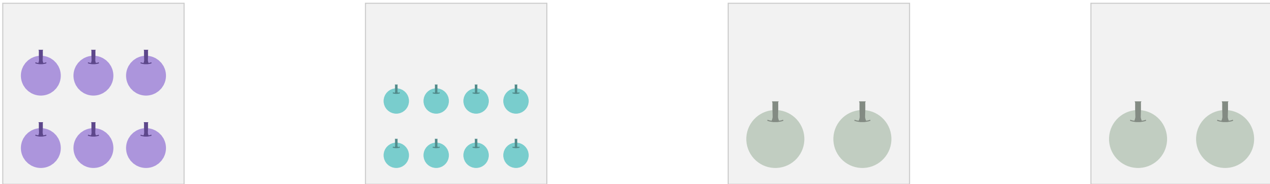

Then we'd need to pick the first $k-1$ baskets and half of the remaining $n-k+1$; this requires us to have

$$(k-1) + \left\lceil \frac{n-k+1}{2} \right\rceil$$

baskets. So we'll try to prove a similar upper bound.

Theorem 3.2

It is always possible to use at most $k + \lfloor \frac{1}{2}(n-k) \rfloor$ baskets.

(These quantities match when $n \not\equiv k \pmod 2$; when $n \equiv k \pmod 2$ they're off by 1.)

Proof. We'll consider n points in $\mathbb{R}^k$ in general position (i.e., such that no hyperplane contains more than k of these points). We'll assign one basket to each point. Then this gives us k measures in $\mathbb{R}^k$, one for each fruit — to obtain the measure corresponding to a certain fruit, we assign each point the number (or fraction) of this fruit in its corresponding basket. For example, if we have one apple in the first basket and two apples in the second, then for the apple measure we'd put a weight of 1 at the point corresponding to the first basket, and a weight of 2 at the point corresponding to the second.

We'll then use the ham sandwich theorem.

Theorem 3.3 (Ham sandwich)

Given k finite measures in $\mathbb{R}^k$, there exists a hyperplane $\mathcal{H}$ such that $\mu_i(\mathcal{H}^+) \geq \frac{1}{2}\mu_i(\mathbb{R}^k)$ and $\mu_i(\mathcal{H}^-) \geq \frac{1}{2}\mu_i(\mathbb{R}^k)$ for each measure μ_i, where $\mathcal{H}^+$ and $\mathcal{H}^-$ are the two *closed* half-spaces created by $\mathcal{H}$.

(By *finite* we mean that $\mu(\mathbb{R}^k) < \infty$.)

Remark 3.4. The ham sandwich theorem is often written for measures which are absolutely continuous with respect to the Lebesgue measure. In this case we can get *exactly* half — we can find $\mathcal{H}$ such that $\mu_i(\mathcal{H}^+) = \mu_i(\mathcal{H}^-) = \frac{1}{2}\mu_i(\mathbb{R}^k)$.

Now we'll apply the ham sandwich theorem to our k measures. Since we chose our points to be in general position, $\mathcal{H}$ itself contains at most k points. This means one of the closed half-spaces $\mathcal{H}^+$ and $\mathcal{H}^-$ has at most $k + \lfloor \frac{1}{2}(n-k) \rfloor$ points, and by the result of the ham sandwich theorem, the baskets in this half-space have at least half of every type of fruit. □

Remark 3.5. This problem also has non-topological solutions (which are nice for $k = 2$), and has appeared in math olympiads.

§4 The ham sandwich theorem

Next, we'll discuss the ham sandwich theorem. We'll prove the version for absolutely continuous measures (for our proof, it's actually enough just that every hyperplane has measure 0).

Theorem 4.1 (Ham sandwich, 1938)

Let $\mu_1, \ldots, \mu_k$ be measures in $\mathbb{R}^k$ which are finite and absolutely continuous (with respect to the Lebesgue measure). Then there exists a 'halving hyperplane' for each of these measures — i.e., a hyperplane $\mathcal{H}$ such that $\mu_i(\mathcal{H}^+) = \mu_i(\mathcal{H}^-) = \frac{1}{2}\mu_i(\mathbb{R}^k)$ for each $1 \leq i \leq k$.

Remark 4.2. The author imagined splitting a leg of ham between two people such that they got the same amount of meat, bone, and so on; this is where the name evolved from.

Proof. As before, we'll want to parametrize the space of objects we're dealing with (here, 'candidate hyperplanes') in some way; this parametrization is again going to be a sphere.

We'll think of μ_k as 'special.' Now consider $\mathbb{S}^{k-1} \subseteq \mathbb{R}^k$. Given any direction $v \in \mathbb{S}^{k-1}$, we can consider a hyperplane perpendicular to v, and we can imagine sliding it along the direction of v until the moment it splits μ_k into two. (It's possible there's an entire range for which the hyperplane splits μ_k in two, if μ_k has no weight on some slice. In that case, we choose the hyperplane in the middle of the range.) We call this hyperplane $\mathcal{H}_v$ — so $\mathcal{H}_v$ is a halving hyperplane for μ_k which is orthogonal to v. We let A be the half-space on the side of $\mathcal{H}$ corresponding to v, and B the half-space on the side corresponding to $-v$. All hyperplanes which split μ_k in half can be parametrized in this way.

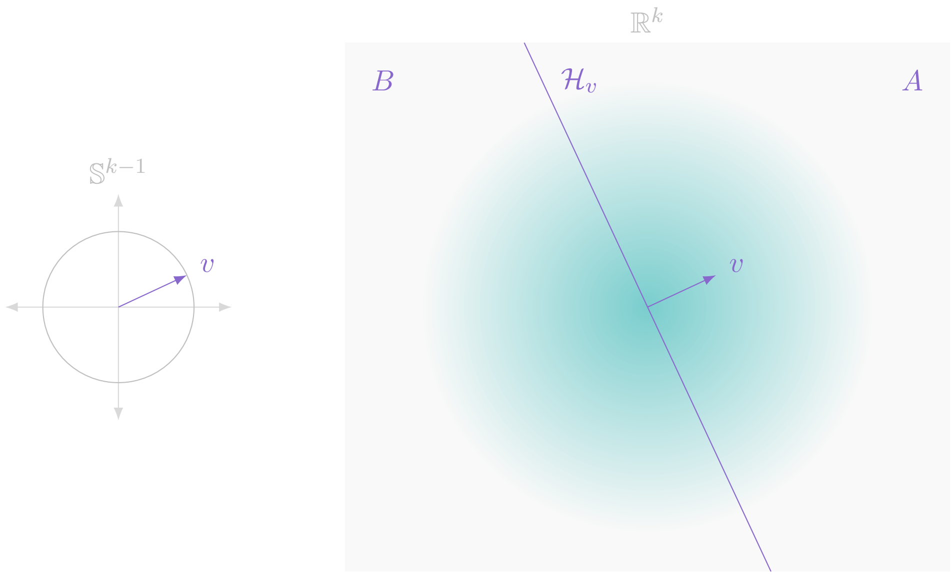

Now we want to construct a function $f\colon \mathbb{S}^{k-1} \to \mathbb{R}^{k-1}$; we define this function as

$$f(v) = (\mu_1(A) - \mu_1(B), \ldots, \mu_{k-1}(A) - \mu_{k-1}(B)).$$

This function is odd (i.e., $f(-v) = -f(v)$ for all v, since negating v keeps $\mathcal{H}$ the same but swaps A and B). By Borsuk–Ulam there must exist v such that $f(v) = f(-v)$, which then means $f(v) = 0$. (This is actually the usual presentation of Borsuk–Ulam — that if we have an odd function $f\colon \mathbb{S}^d \to \mathbb{R}^d$, then it must have a zero.) Here this means $\mathcal{H}_v$ is a halving hyperplane for each of $\mu_1, \ldots, \mu_{k-1}$, and by construction it's also a halving hyperplane for our special measure μ_k. $\square$

§5 The test map scheme

These examples showcase a general method we often follow when using topological methods in combinatorics. (The name 'test map scheme' usually refers to problems about splitting measures, but similar ideas are used more generally.) When we're presented with a problem like the ones we've seen here, we can do the following:

(1) First try to find a space X that parametrizes the objects in the problem.

(2) Then try to find a space Y that gives us information about X. (If we're trying to split many measures, then Y might tell us whether we're giving more of a measure to A or B, or how much of the measure we're giving to A. In our proof of colorful Radon, it told us where in our configuration a point on the octahedron was mapped to.) This induces a natural function $f\colon X \to Y$.

(3) We then try to study the maps $f\colon X \to Y$ as a topological problem, and see if we can say something interesting about them. There's some usual things we're looking for. Maybe X and Y have some symmetry — for example, we can flip half-spaces (in the ham sandwich theorem), and if we have two thieves then we can flip who gets which piece (in the necklace splitting problem). If we have a group G acting on both X and Y, then often we can say things about G-equivariant maps (maps $f\colon X \to Y$ such that applying the G-action before vs. after applying f give the same result). (In all of our examples, G was $\{\pm 1\}$, but there's many situations in which we'll have more complicated groups — for example, the necklace-splitting problem with n thieves. In those cases, the topological result analogous to Borsuk–Ulam will be more complicated.)

The next step from Borsuk–Ulam is the following theorem.

Definition 5.1. For a nontrivial group G, we say a topological space X is a *G-space* if G has a nontrivial free action on X — i.e., an action such that $gx \neq x$ for all $x \in X$ and $g \in G \setminus \{e\}$ — such that the map $x \mapsto gx$ is continuous for each $g \in G$.

Theorem 5.2 (Dold 1985)

Let G be a nontrivial group, and suppose that X and Y are G-spaces which are both paracompact, and such that Y is at most n-dimensional and X is n-connected. Then there is no continuous G-equivariant map $f\colon X \to Y$.

(We won't worry about the meaning of *paracompact*; all our sets will satisfy it. By *n-connected* we mean that the first n reduced homology groups cancel; there's an equivalent definition requiring that whenever you map to a lower-dimensional sphere, there is no homotopy.)

When $G = \mathbb{Z}_2$, $X = \mathbb{S}^{n+1}$, and $Y = \mathbb{S}^n$, this is exactly the Borsuk–Ulam theorem. Dold's theorem is one of the nicer generalizations of Borsuk–Ulam — we only need to check that X and Y have these properties, and that the action of G is free (it's actually enough to check that it's free on X).

§5.1 Necklace splitting for multiple thieves

We'll now use this to solve the necklace splitting problem with more than two thieves. First, if we have k thieves, what do we expect the answer to be? Again, if all the pearls of the same type are clumped together, then we'll need $(k-1)m$ cuts (since for each type of pearl, we need $k-1$ cuts to split this type into k pieces). So we need at least $(k-1)m$ cuts; and it turns out that this is enough.

Theorem 5.3 (Necklace splitting for multiple thieves)

Suppose that we have m continuous measures $\mu_1, \dots, \mu_m$ on $[0,1]$ (which we think of as a continuous version of pearls of m colors on an open necklace), which we want to split evenly among k thieves. It is always possible to find such a fair partition using at most $(k-1)m$ cuts.

First we'll see another solution for two thieves, because it's nice.

Proof for $k = 2$ thieves. We'll map $[0,1] \to \mathbb{R}^m$ using the *moment curve* — we map $\lambda \mapsto (\lambda, \lambda^2, \dots, \lambda^m)$. So now we have m measures in $\mathbb{R}^m$, and we can apply the ham sandwich theorem. This gives a hyperplane $\mathcal{H}$ such that $\mu_i(\mathcal{H}^+) = \mu_i(\mathcal{H}^-)$ for each measure μ_i. (Here we're using the version of ham sandwich for absolutely continuous measures; our measures aren't actually absolutely continuous, because they're concentrated on a curve, but they satisfy the property that the measure of any hyperplane is 0; this is enough for the ham sandwich theorem to apply.)

Now the cool thing about the moment curve is that a hyperplane can only intersect it in m places. So we make the m cuts in these places, and we give every segment in $\mathcal{H}^+$ to A and every segment in $\mathcal{H}^-$ to B; this gives a fair partition. $\square$

Next, we'll consider the case where we have p thieves, where p is prime. (We'll discuss what happens in the non-prime case later.)

Proof for $k = p$ thieves. We first need to parametrize the space of partitions. Before, when we had two thieves, we parametrized our partitions by considering a cut and an assignment of ± 1 to each piece. Here, we'll instead label each piece with $1, \dots, p$ (again corresponding to which thief gets it).

We'll use the *topological join* operation. Intuitively, given two spaces X and Y, we define their join $X * Y$ by embedding both in very high-dimensional space in general position, and then taking their 'convex hull' (more precisely, the set of points $(1-t)x + ty$).

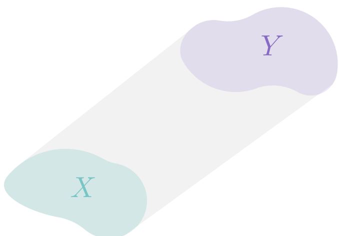

(In this illustration, think of X and Y as being in two separate planes.) Formally, the topological join is defined as follows.

Definition 5.4. For two topological spaces X and Y, we define $X * Y$ as the quotient of $X \times Y \times [0,1]$ by the relation $(x, y, 0) \sim (x, y', 0)$ and $(x, y, 1) \sim (x', y, 1)$ for all $x, x' \in X$ and $y, y' \in Y$.

(In the 'convex hull' interpretation, (x, y, t) corresponds to the point $(1-t)x + ty$ — when $t = 0$ (i.e., we're on the side of X) we don't care what y is, and when $t = 1$ (i.e., we're on the side of Y) we don't care what x is, but there's no other overlap.)

Example 5.5

The space $\mathbb{Z}_2$ consists of two points; then $\mathbb{Z}_2 * \mathbb{Z}_2$ is a quadrilateral; then $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ is an octahedron. More generally, $(\mathbb{Z}_2)^{*n+1} \cong \mathbb{S}^n$ for any n.

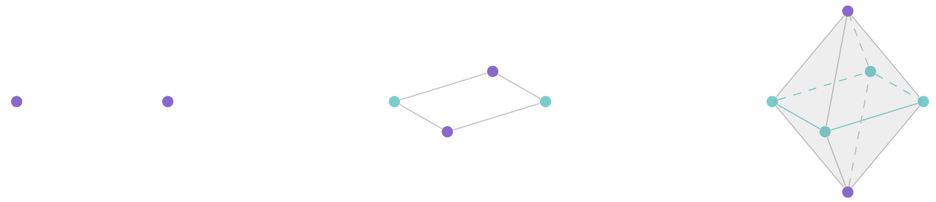

Claim 5.6 — The space of partitions can be parametrized as $(\mathbb{Z}_p)^{*(c+1)}$, where c is our number of cuts.

Proof. Let's first think about $\mathbb{Z}_p * \mathbb{Z}_p$ — this means we have elements $\lambda_1 \in \mathbb{Z}_p$ and $\lambda_2 \in \mathbb{Z}_p$, and some $\alpha \in [0,1]$. We can think of this as the partition $[\alpha \mid 1-\alpha]$, where the first piece (of length α) goes to person λ_1, and the second to person λ_2. If $\alpha = 0$, then we don't care who receives the first piece (since it has length 0), and if $\alpha = 1$, then we similarly don't care who receives the second piece. So $\mathbb{Z}_p * \mathbb{Z}_p$ describes the space of one-cut partitions, and similarly $(\mathbb{Z}_p)^{*(c+1)}$ describes the space of c-cut partitions. □

We want to make $(p-1)m$ cuts, so we consider $(\mathbb{Z}_p)^{*((p-1)m+1)}$. This space is $((p-1)m-1)$-connected — $\mathbb{Z}_p$ is disconnected (i.e., (-1)-connected), and every time we join another $\mathbb{Z}_p$, the connectedness increases by 1. (For example, with $\mathbb{Z}_2$ we started off with a disconnected set, then got a connected set, then a 1-connected set, and so on.)

Now we need to create a space Y and a test map that tell us whether our partition is good. To do so, given a partition, we'll let A_1 be the part that person 1 gets, A_2 the part that person 2 gets, and so on. We also have an action of $\mathbb{Z}_p$ on $(\mathbb{Z}_p)^{*(c+1)}$ (for any c), and this action is free — this is the only reason why we needed p to be prime.

Now for our construction, we can first imagine taking a partition $\mathcal{Q} \in (\mathbb{Z}_p)^{*((p-1)m+1)}$, and tracking how much of each measure each person is getting — so we can define the $m \times p$ matrix

$$f(\mathcal{Q}) = \begin{bmatrix} \mu_1(A_1) & \mu_1(A_2) & \cdots & \mu_1(A_p) \\ \mu_2(A_1) & \mu_2(A_2) & \cdots & \mu_2(A_p) \\ \vdots & \vdots & \ddots & \vdots \\ \mu_m(A_1) & \mu_m(A_2) & \cdots & \mu_m(A_p) \end{bmatrix},$$

which we can view as an element of $\mathbb{R}^{pm}$. We also have an action of $\mathbb{Z}_p$ on the space of such matrices, which exchanges the columns; and this action is equivariant (whether we act before or after applying f, it corresponds to just relabelling who gets each piece).

However, the dimension of this space is too big, so we need to make it smaller. To do so, we'll shift each row to have sum 0 — given a row $[\beta_1, \dots, \beta_p]$, we let $k = \sum_i \beta_i$ be the sum of the row, and we then replace the row with $[\beta_1 - \frac{k}{p}, \dots, \beta_p - \frac{k}{p}]$. Now since each row sums to 0, it can be viewed as an element of $\mathbb{R}^{p-1}$ (embedded into $\mathbb{R}^p$ as the set of vectors with sum of coordinates 0), so our space of matrices can now be viewed as $\mathbb{R}^{(p-1)m}$. So now we have a new function $g\colon (\mathbb{Z}_p)^{*(p-1)m+1} \to \mathbb{R}^{(p-1)m}$.

And a fair partition is precisely a zero of g, so we want to show there exists a zero. Assume not, and define

$$h(\mathcal{Q}) = \frac{g(\mathcal{Q})}{\|g(\mathcal{Q})\|}$$

(this is defined for every partition $\mathcal{Q}$, since we assumed $g(\mathcal{Q})$ is never 0). Then h is a map $(\mathbb{Z}_p)^{*(p-1)m+1} \to \mathbb{S}^{(p-1)m-1}$ (which is continuous and $\mathbb{Z}_p$-equivariant), so we can apply Dold's theorem to get a contradiction, and we're done. □

Finally, we've solved the problem when k is prime; what happens when it's not prime? There are lots of problems that use similar methods, and often they get solved first for 2, then primes, then prime poewrs, and the general case is much more difficult. However, this problem is an exception — we actually don't need any more topology than what we've already seen for the general case, thanks to the following reduction.

Claim 5.7 — If we can solve the problem for a thieves (i.e., find a fair partition of any necklace using $(a-1)m$ cuts) and for b thieves, then we can solve it for ab thieves.

Proof. Clump our thieves into b groups of a thieves each. We first take our original necklace $[0,1]$, and split it fairly among the b *groups* of thieves — this takes $(b-1)m$ cuts.

Then each group takes the intervals assigned to them, puts these intervals back to back to form a new (smaller) necklace, and makes cuts to split it fairly among the a thieves in the group — this takes $(a-1)m$ cuts per group.

So in total, we need $(b-1)m + b(a-1)m = (ab-1)m$ cuts. □

Remark 5.8. Interestingly, the proof of necklace splitting we saw here is different from Alon's proof from 1987. He did use topological methods, but instead of using results that guarantee *no* equivariant maps (like we did), he used results (with different parameters) that guarantee that there *are* equivariant maps, and he used one of those to obtain the fair partition.

§5.2 Method complexity

There's several topological results used in combinatorics, with varying degrees of complexity. The starting point is the Borsuk–Ulam theorem; then there's Dold's theorem, and some degree arguments (about maps of spheres). Then we start getting into fancier things, like characteristic classes, index theories, spectral sequences, and other things. What we've seen here is much closer to the beginning. There's a nice book on the methods on the 'left' side of this line; it's often nice to be on this side because for lots of these problems, in computational geometry we want *algorithms*, and it's easier to get algorithms out of the Borsuk–Ulam proofs than the more fancy ones.

One of the standard proofs of Borsuk–Ulam uses degree, so why did we put degree further down on this list? Here's an example of what we mean by degree.

In the ham sandwich theorem, we have d measures in $\mathbb{R}^d$, and we want a halving hyperplane. What if instead of d measures, we had $d+1$? Usually with $d+1$ measures, we *can't* find a halving hyperplane — for example, if $d = 2$, we could have three measures concentrated at separate vertices of a triangle. However, it *is* true that we can find a halving *disk*.

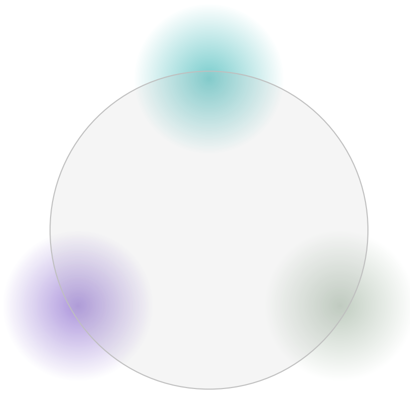

The idea of the proof is to add on a coordinate — for example, we map $(x, y) \mapsto (x, y, x^2 + y^2)$. Now we have $d + 1$ measures in $\mathbb{R}^{d+1}$, and we can apply Ham Sandwich there and project back.

But what if instead of a disk, we want a regular hypercube? For example, when $d = 2$, we have 3 measures and we want to find a square containing exactly half of each.

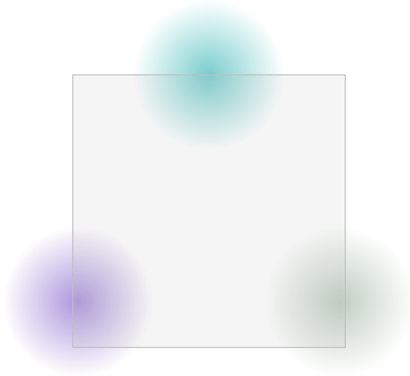

This is still true, but applying standard Borsuk–Ulam tricks doesn't work; the proof is more tricky (even though the topological methods involved may not be more complicated). We'll prove it for $d = 2$ (i.e., for squares); the general case is harder.

Theorem 5.9

Let μ_1, μ_2, and μ_3 be finite, absolutely continuous measures in $\mathbb{R}^2$. Then we can find a square that contains exactly half of each measure.

Proof. We'll parametrize the set of candidate squares by $\mathbb{S}^1 \times [0, 1]$ in the following way — suppose we're given a direction $v \in \mathbb{S}^1$ and some $\alpha \in [0, 1]$. Then we start by scaling v to the point $\frac{\alpha}{1-\alpha} v$. We take a square centered at this point oriented perpendicular to v, and we blow it up until it contains exactly half of μ_3.

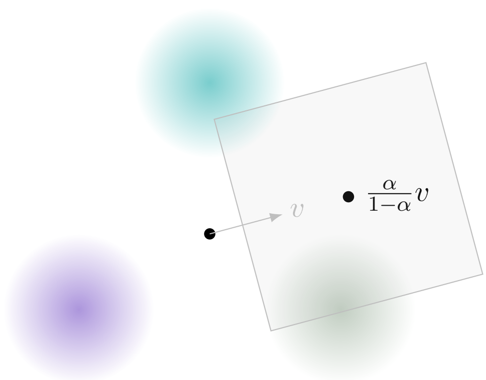

(In particular, we're only considering squares oriented towards the origin.)

By scaling, we can assume our measures are probability measures (i.e., $\mu_i(\mathbb{R}) = 1$ for each i). Then for any candidate square $\mathcal{C}$, we define

$$f(\mathcal{C}) = \left(\mu_1(\mathcal{C}) - \frac{1}{2}, \mu_2(\mathcal{C}) - \frac{1}{2}\right).$$

(So f checks whether our square has half of μ_1 and μ_2.) It suffices to show that f has a zero; assume not, and as before, define $g\colon \mathbb{S}^1 \times [0, 1] \to \mathbb{S}^1$ as

$$g(\mathcal{C}) = \frac{f(\mathcal{C})}{\|f(\mathcal{C})\|}.$$

Now we'll consider what happens when our element of $[0,1]$ is 0 or 1. At 0 we have $g(v,0) = g(-v,0)$ — at 0 we're taking the square centered at the origin and oriented perpendicular to v, and flipping the direction of v doesn't affect this square.

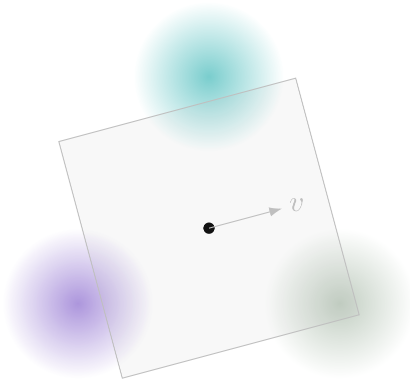

Meanwhile, at 1 we have $g(v,1) = -g(-v,1)$. To see this, as $a \to 1$ the center of our square grows further and further away, and our square keeps growing and growing; this means we're approaching the halving hyperplane for μ_3 perpendicular to v. And if we do this with $-v$ instead, then we'll approach the same halving hyperplane, but the interior of our square will be on the other side. So if we call the two sides of this hyperplane A and B, then we have $\mu_i(A) + \mu_i(B) = 1$, and therefore $\mu_i(A) - \frac{1}{2} = -(\mu(B) - \frac{1}{2})$ (for $i = 1$ and 2). This means $f(v,1) = -f(-v,1)$, so the same is true of g.

But then $g(\,\cdot\,,0)$ is a function $\mathbb{S}^1 \to \mathbb{S}^1$ satisfying $g(v,0) = g(-v,0)$, so it must have even degree; and $g(\,\cdot\,,1)$ satisfies $g(v,1) = -g(-v,1)$, so it must have odd degree. This means g is a homotopy between an even-degree and an odd-degree function; this cannot happen. □

In general, it's not that degree methods are more advanced, but proofs using degree are often trickier.

§6 A linear algebra proof of colorful Radon

Finally, we'll see a solution to the first problem without topology. To restate the problem:

Theorem 6.1 (Colorful Radon, Lovász 1992)
Imagine that we are given $d+1$ pairs of points in $\mathbb{R}^d$, where each pair is of a different color. Then there is a colorful partition — a partition of the points into two sets such that each part contains one point of each color — such that the convex hulls of the two parts intersect.

Proof. Let the pairs of points be $(x_1, y_1), \dots, (x_{d+1}, y_{d+1})$. Consider the vectors $x_i - y_i$ for all $1 \le i \le d+1$. This gives us $d+1$ vectors in $\mathbb{R}^d$, so they must be linearly dependent — so there exist $\alpha_1, \dots, \alpha_{d+1}$ (not all zero) such that

$$\sum \alpha_i (x_i - y_i) = 0.$$

Now, the trick in the proof is that we can assume without loss of generality the α_i are all nonnegative — if some α_i is negative, then we can just swap the names of x_i and y_i (and flip the sign of α_i). And then since they're nonnegative and not all zero, we can also assume that $\sum \alpha_i = 1$ (by scaling). Then by rearranging our equality we get

$$\sum \alpha_i x_i = \sum \alpha_i y_i,$$

which gives a point in the convex hull of both the x_i's and the y_i's. □