Borsuk's Conjecture and
The Chromatic Number of Euclidean Space

Hugo Eberhard

August 2023

1 Introduction

This lecture aims to show how the construction of a graph, partly using some tools from linear algebra, leads to the solution of two seemingly unrelated and very old problems. The first of the problems is called **Borsuk's conjecture**, and the second one is a generalisation of the **Hadwiger–Nelson problem**.

The lecture will closely follow some lectures given by Béla Bollobás on these problems as part of the Part III Combinatorics course at University of Cambridge in Michaelmas in 2022.

1.1 Borsuk's Conjecture

First, let's clear up some terminology. The **diameter** of a set is the distance between the two points in the set that are furthest apart (technically, it's the supremum of all possible distances between points in the set). Note that for a circle or a sphere, this coincides with the usual definition.

In 1932, Borsuk showed that a ball (that is, a solid sphere) can be dissected into four parts, each of which has diameter strictly smaller than the diameter of the ball. In fact, he showed something more general: an $n-$dimensional ball can be covered with $n + 1$ sets, each of which has diameter smaller than the ball. This led him to conjecture the following:

Conjecture 1 ***(Borsuk's Conjecture)*** *Any bounded subset of* $\mathbb{R}^n$ *can be split into* $n+1$ *sets, each of which has smaller diameter than the original set.*

For a long time, people thought the conjecture was true, but no one managed to find a proof. Progress was made several times, in 1946 it was shown to be true for smooth convex sets, in 1971 it was shown to be true for centrally-symmetric sets and in 1995 for bodies of revolution. It was therefore a big surprise when in 1993 Kalai and Kahn showed that the conjecture was *false*. Not only did the conjecture turn out to be false, but it was *very false*, in the sense that the largest number of sets of diameter smaller than 1 you might need to cover a set of diameter 1 in n dimensions, grows exponentially in n. The proof is surprisingly simple, and by the end of this lecture we will have an idea of how it goes.

1.2 The Chromatic Number of Euclidean Space

How many different colours do you need to colour the plane in such a way that no points that are exactly distance one apart have the same colour? It's not that hard to show that seven colours are enough, for example by colouring the plane as in the picture below.

On the other hand, three colours is not enough, which can be seen by trying to colour the vertices of the black triangles in the same picture - you will need at least four colours to make sure no points at distance 1 have the same colour.

The exact answer is actually still unknown, it has so far been narrowed down to either $5, 6$ or 7. The fact that 4 colours is also not enough wasn't proven until 2018 when someone found a set of 1581

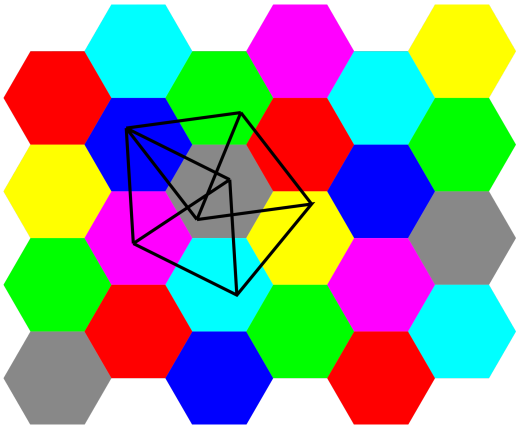

points in the plane which can not be coloured with 4 colours.

So we're sort of stuck trying to find the exact number of colours needed for the plane. But what about higher dimensional spaces? Intuitively it should be harder to pin down the answer, and it is, but what if we instead ask for an approximate answer, or even just a lower bound for the number of colours needed? The following was conjectured by Erdős at some point (I don't know when):

Conjecture 2 *The number of colours needed to colour $\mathbb{R}^n$ such that no points at distance 1 have the same colour, is exponential in n.*

For many years, no one managed to prove this, but by the end of this lecture we will have seen a simple proof that the conjecture was true.

2 $L-$intersecting families

In the combinatorics lecture, we talked about some classic problems, one of which was the problem of Oddtown. We showed that in a town with n people who are members of some different clubs, assuming that each club has an odd number of members and that any two clubs have an even number of members in common, there can be at most n clubs in total. We can generalise this notion, and talk about $L-$*towns*.

Definition 1 *Let L be a set of (non-negative) integers. If we have a town with n people who are members of some different clubs, we say that the town is an $L-$town if the following holds:*

- *The number of members in a club is never in the set L*
- *The number of members that any two clubs have in common is in the set L*

We tend to think of the people in the town as the set $[n] = \{1, 2, ..., n\}$, and the clubs as subsets $A_1, ..., A_m$ of this set. Then an $L-$town is exactly a family of $L-$intersecting subsets, specifically we require that

- *$|A_i| \notin L$ for all i*
- *$|A_i \cap A_j| \in L$ for all $i \neq j$*

We can similarly talk about $L-$towns modulo p, for any prime p. The definition is exactly the same, except that the conditions on $|A_i|$ and $|A_i \cap A_j|$ are considered modulo p.

Our first aim is to, given $|L|$ and n, prove a bound on the number of clubs in an $L-$town. Equivalently, we want to bound the size of an $L-$intersecting family. We have the following theorem:

Theorem 1 *Let $L = \{l_1 < ... < l_s\}$ be a set of s integers, and let $A_1, ..., A_m \subset [n]$ be an*

$L-$intersecting family of m sets. Then the following holds:

$$m \leq \binom{n}{0} + \binom{n}{1} + ... + \binom{n}{s}$$

The result holds regardless of whether we consider an $L-$intersecting family modulo p or not.

Let's check that this agrees with our previous results. If we consider $L-$intersecting families modulo $p = 2$ and $s = |L| = 1$ the theorem says $m \leq n + 1$. This exactly corresponds to the example with Oddtown (if $L = \{0\}$), and the theorem gives a bound which is very close to our previous bound.

Proof of theorem: As usual, we think of the sets (or clubs) $A_1, ..., A_m \subset [n]$ as $n-$dimensional indicator vectors, say $v_1, ..., v_m$ where v_i has a 1 in position j if and only if $j \in A_i$. We will work in $\mathbb{R}^n$ in the case where we don't consider the conditions modulo p, and in $\mathbb{F}_p^n$ when we consider the conditions modulo p. Both $\mathbb{R}$ and $\mathbb{F}_p$ are fields, so linear algebra works as usual (although in the second case, note that dot products are not technically inner products, since $\langle v, v \rangle = 0 \pmod p$ does not imply $v = 0$).

As usual, we define the dot product as

$$\langle v, w \rangle = v_1 w_1 + v_2 w_2 + ... + v_n w_n$$

The key of the proof will be to consider the following polynomials. For each vector v_i corresponding to a set A_i we define:

$$f_i(x_1, ..., x_n) = (\langle v_i, x \rangle - l_1)(\langle v_i, x \rangle - l_2)...(\langle v_i, x \rangle - l_s)$$

This polynomial has degree s and is 0 if and only if $\langle v_i, x \rangle$ is in L.

We define associated polynomials $\tilde{f}_i(x)$ by expanding the factors of $f_i(x)$ and replacing every occurrence of x_j^k for some k by x_j. For example, if

$$f_i(x_1, x_2, x_3) = 5x_1 x_2^3 x_3 - x_2 x_3^2 + x_1^5 + 4x_2^2 x_3^3$$

we get

$$\tilde{f}_i(x_1, x_2, x_3) = 5x_1 x_2 x_3 - x_2 x_3 + x_1 + 4x_2 x_3 = 5x_1 x_2 x_3 + 3x_2 x_3 + x_1$$

Note that if we only care about vectors x such that all entries x_j are 0 or 1, the value of x_j^k is the same as x_j, and hence the following also holds for $\tilde{f}_i(x)$:

- For any x with entries in $\{0, 1\}$, we have that $\tilde{f}_i(x) = 0$ if and only if $\langle v_i, x \rangle \in L$. In particular, since the v_j have entries only in $\{0, 1\}$ and $\langle v_i, v_j \rangle = |A_i \cap A_j|$ we get that for any $L-intersecting$ family, $\tilde{f}_i(v_j) = 0$ if and only if $i = j$.
- $\tilde{f}_i(x)$ has degree at most s

The first of these observations implies that the $\tilde{f}_i$ are linearly independent (seen as elements of the vector space of polynomials in $x_1, ..., x_n$ over $\mathbb{R}^n$ or $\mathbb{F}_p^n$). Indeed,

$$\sum \lambda_i \tilde{f}_i(x) = 0 \quad \Longrightarrow \quad \sum \lambda_i \tilde{f}_i(v_j) = 0 \quad \Longrightarrow \quad \lambda_j \tilde{f}_j(v_j) = 0 \quad \Longrightarrow \quad \lambda_j = 0$$

The second observation means that the polynomials are multilinear polynomials of degree at most s, that is they live in the span of the polynomials of form $x_{i_1} x_{i_2} ... x_{i_k}$ where $0 \leq k \leq s$. There are exactly $\binom{n}{0} + \binom{n}{1} + ... + \binom{n}{s}$ such polynomials, so the dimension of this vector space is at most (in fact exactly)

$$\binom{n}{0} + \binom{n}{1} + ... + \binom{n}{s}$$

Hence we are done using that $\tilde{f}_i$ are all linearly independent. ■

It's worth noting that the version of the theorem which does not consider sizes modulo p follows immediately from the modulo p version: just consider a prime $p > n$. However the opposite implication is less clear.

This is already a pretty strong theorem. For example, considering $L = \{0, 1, 2, ..., s-1\}$, we note that the family of all subset of size s has size $\binom{n}{s}$ and is $L-$intersecting (in the modulo p case we also need $p > s$ for this to work). Hence we have an upper bound on the number of sets (or clubs) m which is a polynomial in n of the same degree as the lower bound. However, to make progress on the two problems from the introduction, we will need a different but very similar theorem:

Theorem 2 *Let r be an integer and let $L = \{l_1 < ... < l_s\} \subset \{0, 1, ..., r-1\}$ be a set of s integers. Let $A_1, ..., A_m \subset [n]$ be an $L-$intersecting family of m sets. If it additionally holds that $|A_i| = r$ for all i, and that $r + s \leq n$, then the following holds:*

$$m \leq \binom{n}{s}$$

This theorem is in fact also true in the modulo p case, but the proof of that is a bit trickier than what we will show here. We will however prove a special case of the modulo p version of the theorem later.

Proof of theorem: We work in $\mathbb{R}[x_1, ..., x_n]$, the set of real polynomials in $x_1, ..., x_n$, and view this as a vector space over $\mathbb{R}$. As usual, think of the sets $A_1, ..., A_m$ as indicator vectors $v_1, ..., v_m$ in $\{0, 1\}^n \subset \mathbb{R}^n$. We make the following definitions:

- For each set $I = \{i_1, ..., i_k\} \subset [n]$ let $m_I(x) = x_{i_1}x_{i_2}...x_{i_k}$.
- Let $p_I(x) = m_I(x)(\sum_{i=1}^{n} x_i - r)$.
- As before, define for each set A_i an associated polynomial

$$f_i(x) = (\langle v_i, x\rangle - l_1)(\langle v_i, x\rangle - l_2)...(\langle v_i, x\rangle - l_s)$$

- For any polynomial $g(x)$, define $\tilde{g}(x)$ in the same way as in the last proof (i.e. we obtain $\tilde{g}$ by replacing all occurrences of x_j^k for some $k > 1$ by x_j)

As before, we know that $\tilde{f}_i(x)$ are linearly independent and that they are all in the span of

$$M = \{m_I \mid |I| \leq s\}$$

which is a subspace of dimension $\binom{n}{0} + ... + \binom{n}{s}$. Furthermore, we claim that the set

$$\{\tilde{f}_1, ..., \tilde{f}_m\} \cup \{\tilde{p_I} \mid |I| \leq s-1\}$$

is linearly independent. Note that it has size $m + \binom{n}{0} + ... + \binom{n}{s-1}$ and that all the polynomials are in the span of M, so if we can show this we are done.

First we show that the $\tilde{p_I}(x)$ are all linearly independent. Indeed, assume $\sum \lambda_I \tilde{p_I}(x) = 0$. Evaluating at the indicator vector v_J of a set $J \subset [n]$ gives that

$$\tilde{p_I}(v_J) = p_I(v_J) = m_I(v_J)(|J| - r) = \begin{cases} 0 & \text{if } |J| < |I| \text{ or if } |I| = |J| \text{ and } I \neq J \\ |I| - r \neq 0 & \text{if } J = I \end{cases}$$

Hence if we evaluate the sum $\sum \lambda_I \tilde{p_I}(x)$ at v_J for all $J \subset [n]$ or size at most $s-1$ *in the order from smallest to largest*, at any point of the process the sum will simplify to just $\lambda_J(|J| - r) = 0$, giving that $\lambda_J = 0$ since $|J| \leq s - 1 < r$ (note that the order is important, since we can't say anything about $p_I(v_J)$ for sets I such that $|I| < |J|$, but we will already know that $\lambda_I = 0$ for such I if we do it in the correct order). This shows that the $\tilde{p_I}(x)$ are linearly independent.

Finally, we show that the polynomials $\tilde{f_i}$ and $\tilde{p_I}$ are linearly independent. Assume that $\sum \lambda_i \tilde{f_i} = \sum \mu_I \tilde{p_I}$. Evaluating at v_j (the indicator vector of the j^{th} set) we get that $\tilde{p_I}(v_j) = p_I(v_j) = m_I(v_j)(r - r) = 0$ using that $|A_j| = r$, while $\tilde{f_i}(v_j) = f_i(v_j) = 0$ unless $i = j$ (as $A_1, ..., A_m$ is an $L-$intersecting family, the reasoning here is the same as in the previous proof). We conclude that $\lambda_j = 0$, for every j, and so $\sum \mu_I \tilde{p_I} = 0$ giving that $\mu_I = 0$ for every I as well (as the $\tilde{p_I}$ are linearly independent). Hence we are done. ■

Reading this proof carefully, we see that every part of it works over $\mathbb{F}_p^n$ as well, *except* when we proved that the polynomials $\tilde{p_I}$ for $|I| \leq s-1$ are linearly independent. The reason this step might not work is that $p_I(v_I) = |I| - r$ might not be non-zero over $\mathbb{F}_p$. In the proof above, we had $|I| \le s - 1 < r$ which was enough to conclude it's non-zero, but in the modulo p case we need to be more careful. However, if we let $r = 2p - 1$ and $s = p - 1$, it's clear that

$$|I| - r = |I| - 2p + 1 \in \{1 - 2p,\ 2 - 2p,\ \ ...,\ (p-1) - 2p\}$$

in all the cases we care about. Crucially, it's not 0 modulo p, so the proof carries over exactly as written. We hence get the following corollary by picking $L = \{0, 1, 2, ..., p - 2\}$:

Corollary 1 *Let p be a prime and let $n = 4p - 1$. If $A_1, ..., A_m$ are subsets of $[n]$ of size $2p - 1$ such that $|A_i \cap A_j| \neq p - 1$, then*

$$m \le \binom{4p-1}{p-1}$$

3 Borsuk's Conjecture and The Chromatic Number of Euclidean Space

We now have the tools to tackle the problems from the introduction. Both the proofs that we are about to see rely on considering $n = 4p - 1$ for some prime p and constructing a graph G as follows:

- Let the vertex set $V(G)$ be the set of all subsets of $[n]$ of size $2p - 1$. We can view this a subset of $\{0, 1\}^n \subset \mathbb{R}^n$, where a set A is associated with the point $x \in \{0, 1\}^n$ such that $x_i = 1$ if and only if $i \in A$.
- Put an edge between A and B if and only if $|A \cap B| = p - 1$.

The key point will be that for any vertices A and B that are connected by an edge, the distance between the corresponding points in $\mathbb{R}^n$ is $\sqrt{2p}$. In particular, all edges have the same length (if we embed the graph in $\mathbb{R}^n$ as described above).

Theorem 3 *For large enough n, the number of colours needed to colour $\mathbb{R}^n$ such that no two points at distance 1 have the same colour is at least 1.05^n.*

Proof: Start by considering n of the form $4p - 1$ for some prime p and define the graph G as above. We embed it in $\mathbb{R}^n$ in the same way as described before. Let us colour the points of $\mathbb{R}^n$ such that no two points at distance $\sqrt{2p}$ have the same colour (finding a lower bound for the number of colours used in this colouring is the same as our original problem, after rescaling).

Consider a set of vertices in the graph that all have the same colour. No two of them can be adjacent in the graph since all edges in the graph have length $\sqrt{2p}$, which means that the sets $A_1, ..., A_m$ corresponding to the vertices are such that $|A_i \cap A_j| \neq p - 1$. By the last corollary from the previous section, this means that $m \le \binom{4p-1}{p-1}$. Hence there can be at most this many vertices in the graph of any given colour, so the number of colours is at least

$$\frac{\binom{4p-1}{2p-1}}{\binom{4p-1}{p-1}} = \frac{3p(3p-1)...(2p+1)}{(2p-1)(2p-2)...p} \ge \left(\frac{3}{2}\right)^p$$

Finally, it's well-known that there is a prime between x and $2x$ for any integer x. Hence for any n we can find a prime $p \ge \frac{n}{8}$ such that $4p - 1 < n$, giving the following lower bound on the number of colours for any n:

$$\#\text{colours} \ge \left(\frac{3}{2}\right)^p \ge \left(\frac{3}{2}\right)^{n/8} > 1.05^n$$

■

Theorem 4 *For large enough N, there exists a subset of $\mathbb{R}^N$ of diameter 1 such that the number of sets of diameter smaller than 1 needed to cover it is at least $1.05^{\sqrt{2N}}$.*

Proof of theorem: Again, start by considering n of the form $4p - 1$ for some prime p and define the graph G as above. Let $N = \binom{n}{2}$. This time, we will embed the graph in $\mathbb{R}^N = \mathbb{R}^{\binom{n}{2}}$ in the following way.

Associate the vertex A with the set V_A defined by

$$V_A = \{\{x, y\} \mid x \in A, y \notin A\}$$

Note that for subsets A, B of $[n]$ of size $2p-1$,

$$\begin{aligned} |V_A \setminus V_B| + |V_B \setminus V_A| &= 2\big(|A \cap B|(|B| - |A \cap B|) + (|A| - |A \cap B|)(n - |A| - |B| + |A \cap B|)\big) \\ &= 2(2p - 1 - |A \cap B|)(1 + 2|A \cap B|) \end{aligned} \tag{1}$$

This is maximized when $|A \cap B| = p - \frac{3}{4}$, so since $|A \cap B|$ is an integer the unique maximum is $|A \cap B| = p - 1$, so exactly when there is an edge between A and B in the graph G.

Now, the sets V_A can be associated with points in $\{0, 1\}^N \subset \mathbb{R}^N$ in the usual way (note that V_A is a subset of the subsets of $[n]$ of size 2, which is a set of size $N = \binom{n}{2}$). This let's us embed G into $\mathbb{R}^N$ in such a way that the distance between any two vertices is given by the formula in equation (1). Note in particular that any two vertices connected by an edge are at distance $\sqrt{2p(2p-1)}$ from each other, and that all other pairs of vertices are closer to each other. Hence, partitioning the set of vertices into sets of diameter $< \sqrt{2p(2p-1)}$ exactly corresponds to finding a colouring of the vertices of G in which no adjacent vertices have the same colour. By the same reasoning as in the previous proof, we get that we need at least $\left(\frac{3}{2}\right)^p$ sets. Using that $n > \sqrt{2\binom{n}{2}} = \sqrt{2N}$ and that for any n there is a prime $p \geq \frac{n}{8}$ such that $4p - 1 < n$, we get a lower bound of $1.05^{\sqrt{2N}}$ in general. ■