Week 16

Tibor Szabó

Extremal Combinatorics, FU Berlin, WiSe 2017–18

1 The Borsuk–Ulam Theorem

In this section we shall see another kind of application of the topological method in combinatorics — we shall use a purely topological theorem, the Borsuk–Ulam Theorem, to prove Kneser's Conjecture, an entirely combinatorial result. This is a famous result, having given birth to topological combinatorics, and we begin by introducing the conjecture.

1.1 Kneser's Conjecture

Recall the Erdős–Ko–Rado Theorem: when $n \geq 2k$, the size of a k-uniform intersecting family of subsets of $[n]$ is at most $\binom{n-1}{k-1}$. (a family is intersecting if any two members have a non-empty intersection). This is tight: any *star*, the family of all k-subsets containing a fixed element, has size exactly $\binom{n-1}{k-1}$.

Let us reformulate this in terms of an important classic graph, the Kneser graph $KG(n,k)$.

Definition 1.1 (Kneser graph). *Given integers $n \geq k \geq 1$, the* Kneser graph *$KG(n,k)$ has $V(KG(n,k)) = \binom{[n]}{k}$, the set of all k-sets of n, with $E(KG(n,k)) = \{\{A,B\} : A \cap B = \emptyset\}$.*

For a couple of quick examples, observe that if $n < 2k$, then there are no disjoint pairs, so $KG(n,k)$ consists of $\binom{n}{k}$ isolated matchings. If $n = 2k$, then we have a perfect matching on $\binom{2k}{k}$ vertices, with complementary pairs connected. These graphs are not new to us either — $KG(5,2)$ is better known as the Petersen graph.

What does the Erdős–Ko–Rado Theorem tell us about the Kneser graph? A family $\mathcal{F} \subseteq \binom{[n]}{k}$ is intersecting if and only if it forms an independent set in $KG_{n,k}$, and therefore the theorem says that when $n \geq 2k$, $\alpha(KG(n,k)) = \binom{n-1}{k-1}$.

What did we always do when we studied the independence number of some nice graph or hypergraph? We also looked at the chromatic number, which is related through the fundamental inequality $\chi(G) \geq v(G)/\alpha(G)$. For example, when $n \geq 2k$, this shows the chromatic number of $KG(n,k)$ is at least $\binom{n}{k}/\binom{n-1}{k-1} = n/k$. How good is this lower bound?

We start by looking back at our examples. When $n < 2k$, the Kneser graph is the empty graph, so its chromatic number is 1. For $n = 2k$, $KG(2k,k)$ is a perfect matching, so its chromatic number is 2. When $n = 2k+1$, then we get the so-called Odd-graph. This is already more complicated, with lots of mysteries. It is $(k+1)$-regular, but what other properties does it have? For example, when $k = 1$, the Odd graph is just the triangle K_3, so its chromatic number is three. When $k = 2$, we get the Petersen graph. This has a 5-cycle, and hence is not bipartite. It is, however, possible to color it with three colors. So its chromatic number is also three. Applying the lower bound via the independence number to the odd graphs, we get that the chromatic number is at least $\lceil n/k \rceil = \lceil (2k+1)/k \rceil = 3$. So far, so good.

However, we have only checked the cases $k \leq 2$. Can we color any Odd-graph with three colors properly? How should we start? We can start, like in case of the Petersen graph, by coloring a maximum independent set with one color: a star, say containg the vertex 1. How to continue? We could use one color on the next star, the sets containing 2 but not 1. And so on. The last star we would need to take is the one having the k-sets containing $n-k+1$. This has only one member,

$\{n-k+1, n-k+2, \ldots, n\}$; all other k-sets have smaller minimum element. However, observe that we could save some colors by collapsing some of the smaller stars into one intersecting family: over any $(2k-1)$-element vertex set, any two k-sets intersect. So we could take all k-subsets of the $(2k-1)$-element set $[3, 2k+1]$ and that will be independent. Thus we can indeed color any odd graph with just three colors.

What can we say for general Kneser graphs? We could start with the same trick, taking stars up to the element $n-2k+1$ and then all k-subsets of the set $[n-2k+2, n]$. This is a proper $(n-2k+2)$-coloring. In 1955 Kneser introduced Kneser graphs and conjectured that this coloring is optimal.

Conjecture 1.2. *For every $n \geq 2k$, we have $\chi(KG_{n,k}) = n-2k+2$*

Note that we have already verified this theorem whenever $n = 2k$ and $2k+1$. The general case turned out to be much more difficult. The first proof was given by László Lovász. The significance of this proof cannot be overestimated, as it initiated the field of topological combinatorics. Within a week of the first proof, a second proof was given by Imre Bárány using the Borsuk–Ulam Theorem and the Gale transform. Maybe the simplest possible argument using an appropriate variant of the Borsuk–Ulam Theorem was given by Josh Greene in 2002, while he was an undergraduate student, and this is the proof we will present here. Later, Matousek gave a proof using only Tucker's Lemma — the combinatorial backbone of the Borsuk–Ulam Theorem. It is somewhat surprising that there is no real topology-free proof of this seemingly entirely combinatorial statement.

In the next section we discuss and motivate the Borsuk–Ulam Theorem. Then we use it to prove Lovász's Theorem.

1.2 The Borsuk-Ulam Theorem

Definition 1.3. *The d-*dimensional *sphere $S^d \subseteq \mathbb{R}^{d+1}$ is given by*

$$S^d = \left\{ \mathbf{x} \in \mathbb{R}^{d+1} : x_0^2 + x_1^2 + \cdots + x_d^2 = 1 \right\}.$$

A pair $\{\mathbf{x}, -\mathbf{x}\} \subseteq S^d$ is called antipodal.

The Borsuk–Ulam Theorem tells us that continuous maps from the d-sphere to d-space must have a pair of antipodal points with the same image.

Theorem 1.4 (Borsuk–Ulam Theorem)**.** *If $f : S^d \to \mathbb{R}^d$ is a continuous map, then there exists some $\mathbf{x}^* \in S^d$ such that $f(\mathbf{x}^*) = f(-\mathbf{x}^*)$.*

Remark 1.5.

1. *This theorem has some real-world implications:*
   (a) *In one dimension, there must be opposite points on the Equator with the same temperature along an equator.*
   (b) *In two dimensions, if we deflate and flatten an inflatable ball, two antipodal points will lie on top of each other.*
   (c) *Still in two dimensions, given a point on the Earth's surface, let d be the distance to the nearest McDonald's, and p be the probability of being hit by a meteorite. Then f : Earth's surface $\to \begin{pmatrix} d \\ p \end{pmatrix}$ is continuous (exercise), and so Borsuk–Ulam says there are diametrically opposite points with the exact same distance to the nearest McDonald's AND the same probability of being hit by a meteorite.*
2. *There are many equivalent formulations of the Borsuk-Ulam Theorem, as well as many different proofs (combinatorial, algebraic topological, etc.). We will see an appropriate equivalent formulation in the next sebsection and some more in the homework.*

While we will not prove the Borsuk–Ulam Theorem in full, we can at least prove the one-dimensional case.

Proof of 1-dimensional case. Create a map $g : [0, \pi] \to \mathbb{R}$ by

$$g(\theta) = f(\mathbf{x}(\theta)) - f(-\mathbf{x}(\theta)), \text{ where } \mathbf{x}(\theta) = \begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix}.$$

g is then a continuous function with $g(\pi) = -g(0)$. By the Intermediate Value Theorem, there must be some $\theta^* \in [0, \pi]$ with $g(\theta^*) = 0$. Thus, if $\mathbf{x}^* = \mathbf{x}(\theta^*)$, $f(\mathbf{x}^*) = f(-\mathbf{x}^*)$. □

1.3 A reformulation

The following theorem of Lusternik and Schnirelmann turns out to be equivalent to Theorem 1.4, and will be more convenient for our application.

Theorem 1.6. *If $S^d = U_0 \cup U_1 \cup \cdots \cup U_d$ is a covering of the sphere with $d+1$ sets, where for $1 \le i \le d$, U_i is either open or closed, then for some $0 \le j \le d$, U_j contains an antipodal pair.*

Remark 1.7.

1. *Note that there is no topological restriction on U_0.*
2. *The topology used here is the subspace topology of $S^d \subseteq \mathbb{R}^{d+1}$.*
3. *Antipodal points achieve the diameter of the sphere. The Theorem implies we cannot split $S^d \subseteq \mathbb{R}^{d+1}$ into $d+1$ "nice" parts and decrease the diameter. This motivates Borsuk's Conjecture, which we saw earlier in the course (in fact, it is rather Borsuk's Question: he never conjectured it ...).*

Proof of Theorem. Suppose the theorem is false. Let $U_0 \cup U_1 \cup \cdots \cup U_d$ be a valid cover w/o any part containing an antipodal pair.

Define $f : S^d \to \mathbb{R}^d$ as

$$f(\mathbf{x}) = (d(\mathbf{x}, U_1), d(\mathbf{x}, U_2), \ldots, d(\mathbf{x}, U_d)),$$

where $d(\mathbf{x}, U_i) = \inf_{\mathbf{y} \in U_i} d(\mathbf{x}, \mathbf{y}) = \inf_{\mathbf{y} \in U_i} \|\mathbf{x} - \mathbf{y}\|_2$ in $\mathbb{R}^{d+1}$. Using the triangle inequality, one can show that f is continuous.

By Borsuk-Ulam, there is some $\mathbf{x}^* \in S^d$ such that $f(\mathbf{x}^*) = f(-\mathbf{x}^*)$. We cannot have $\{\mathbf{x}^*, -\mathbf{x}^*\} \subseteq U_0$, as U_0 has no antipodal pairs.

Therefore, WLOG,

$$\begin{aligned} &\mathbf{x}^* \in U_1. \\ &\qquad \Rightarrow d(\mathbf{x}^*, U_1) = 0 \\ &\qquad \Rightarrow d(-\mathbf{x}^*, U_1) = 0 \\ &\qquad \Rightarrow -\mathbf{x}^* \in \bar{U}_1, \end{aligned}$$

where $\bar{U}_1$ is the closure of U_1. If U_1 was closed, $U_1 = \bar{U}_1$, and so $\{\mathbf{x}^*, -\mathbf{x}^*\} \subseteq U_1$. This would be a contradiction.

Therefore, U_1 is open. Observe that $-U_1 := \{-\mathbf{y} : \mathbf{y} \in U_1\}$ is disjoint from U_1, since U_1 doesn't contain antipodal pairs. So, $S^d \setminus (-U_1)$ is a closed set containing U_1. Thus, $\bar{U}_1 \subseteq S^d \setminus (-U_1)$ (since $\bar{U}_1$ is the intersection of all closed sets containing U_1), and so $-\mathbf{x}^* \in S^d \setminus (-U_1)$, implying $-\mathbf{x}^* \notin -U_1$, and thus $\mathbf{x}^* \notin U_1$. again giving a contradiction. □

1.4 Proof of Kneser's Conjecture

We now prove Conjecture 1.2, using these topological tools.

Theorem 1.8 (Lovász)**.** *When $n \ge 2k$, there is no partition of $\binom{[n]}{k}$ into $n - 2k + 1$ intersecting families.*

Proof (Greene, 2002). We want to show that for any partition $\binom{[n]}{k} = \mathcal{F}_1 \cup \mathcal{F}_2 \cup \cdots \cup \mathcal{F}_{n-2k+1}$, one of the parts $\mathcal{F}_i$ contains a disjoint pair of k-sets.

Recall that we say vectors $\{\mathbf{v}_i : 1 \leq i \leq n\}$ are in general linear position in $\mathbb{R}^{d+1}$ if any subset of at most $d+1$ vectors is linearly independent. Let

$$\mathbf{x}_i = \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|} \in S^d \subseteq \mathbb{R}^{d+1}.$$

We then say $\{\mathbf{x}_i : 1 \leq i \leq n\}$ are in general position in S^d. In particular, any equator of S^d (the intersection of S^d with a hyperplane containing the origin) contains at most d of the points $\mathbf{x}_i$.

Let $\{\mathbf{x}_i : 1 \leq i \leq n\}$ be n points in general position in S^{n-2k+1}. Given $\mathbf{y} \in S^{n-2k+1}$, let $H(\mathbf{y})$ be the *open* hemisphere centered at $\mathbf{y}$.

Observation 1.9. $\mathbf{z} \in H(\mathbf{y}) \Leftrightarrow \|\mathbf{z} - \mathbf{y}\| < \sqrt{2}$.

We can now define a cover of the sphere S^{n-2k+1}. Let $U_i, 1 \leq i \leq n-2k+1$, be defined as

$$U_i = \left\{\mathbf{y} \in S^{n-2k+1} : \exists F \in \mathcal{F}_i \text{ such that } \forall j \in F, \mathbf{x}_j \in H(\mathbf{y})\right\}$$

Let $U_0 = S^{n-2k+1} \setminus \left(\bigcup_{i=1}^{n-2k+1} U_i\right)$.

Claim 1.10. *For $1 \leq i \leq n-2k+1$, U_i is open.*

Proof of Claim. Suppose $\mathbf{y} \in U_i$. Then there is some set $F \in \mathcal{F}_i$ such that $\mathbf{x}_j \in H(\mathbf{y})\ \forall j \in F$, and so $\|\mathbf{x}_j - \mathbf{y}\| < \sqrt{2}\ \forall j \in F$.

Let $\varepsilon = \min_{j \in F}\left(\sqrt{2} - \|\mathbf{x}_j - \mathbf{y}\|\right) > 0$. Now if $\mathbf{y}' \in S^{n-2k+1}$ is such that $\|\mathbf{y} - \mathbf{y}'\| < \varepsilon$, then for each $j \in F$, we have

$$\begin{aligned}\|\mathbf{x}_j - \mathbf{y}'\| &\leq \|\mathbf{x}_j - \mathbf{y}\| + \|\mathbf{y} - \mathbf{y}'\| \\ &< \|\mathbf{x}_j - \mathbf{y}\| + \varepsilon \\ &\leq \sqrt{2},\end{aligned}$$

and so $\mathbf{x}_j \in H(\mathbf{y}')$ for all $j \in F$. Hence $\mathbf{y}' \in U_i$. Thus U_i is open. □

By Theorem 1.6, some U_j, $0 \leq j \leq n-2k+1$, contains an antipodal pair $\{\mathbf{y}, -\mathbf{y}\}$.

Claim 1.11. $j \neq 0$.

Proof. Observe that $\mathbf{z} \in U_0 \iff H(\mathbf{z})$ contains at most $k-1$ points $\mathbf{x}_i$. (If it contained k points, they would form a set and so belong to one of the $\mathcal{F}_i$, making $\mathbf{z} \in U_i$ instead.)

If $\mathbf{y}, -\mathbf{y} \in U_0$, then both of their hemispheres contain $\leq k-1$ points and so at least $n-2k+2$ points $\mathbf{x}_i$ lie on the equator equidistant from $\mathbf{y}$ and $-\mathbf{y}$ (that is, $S^{n-2k+1} \setminus (H(\mathbf{y}) \cup H(-\mathbf{y}))$), contradicting the fact that the points are in general position. □

Therefore, some U_i contains antipodal points $\{\mathbf{y}, -\mathbf{y}\}$ for some i with $1 \leq i \leq n-2k+1$. Thus, $H(\mathbf{y})$ contains $\{\mathbf{x}_j : j \in F\}$ for some $F \in \mathcal{F}_i$. But also, $H(-\mathbf{y})$ contains $\{\mathbf{x}_j : j \in F'\}$ for some $F' \in \mathcal{F}_i$. Since the intersection of these *open* hemispheres is empty, we have that these points $\mathbf{x}_j$ are distinct, and this implies $F \cap F' = \emptyset$. □

2 Problems in Combinatorial Geometry

We close our course by discussing a few classic problems of Erdős from combinatorial geometry.

2.1 The problems

We have encountered already the unit distance graph UD_d of the plane and estimated its chromatic number. Now we investigate how dense can its finite subgraphs be. In other words, among n points in $\mathbb{R}^d$, how many pairs can have distance 1?

For a point set $P \subseteq \mathbb{R}^2$, let $u(P)$ denote the number of pairs of points from P that are of distance 1 from each other. Let $u(n) = \max u(P)$ denote the maximum possible number of pairs at unit distance over all sets of n points in the plane. Erdős posed the problem of determining how $u(n)$ grows with n.

Let us consider a few initial constructions. By taking the vertices of a regular n-gon with sidelength 1, we have $u(n) \geq n$. One does better by placing the points in an $\sqrt{n} \times \sqrt{n}$ grid, again with sidelength 1, which gives $u(n) \geq (2 - o(1))n$. A further improvement can be obtained by shearing the grid into a $\sqrt{n} \times \sqrt{n}$ triangular grid, so that one of the diagonals in each of the previous squares also becomes unit length. We then have $u(n) \geq (3 - o(1))n$.

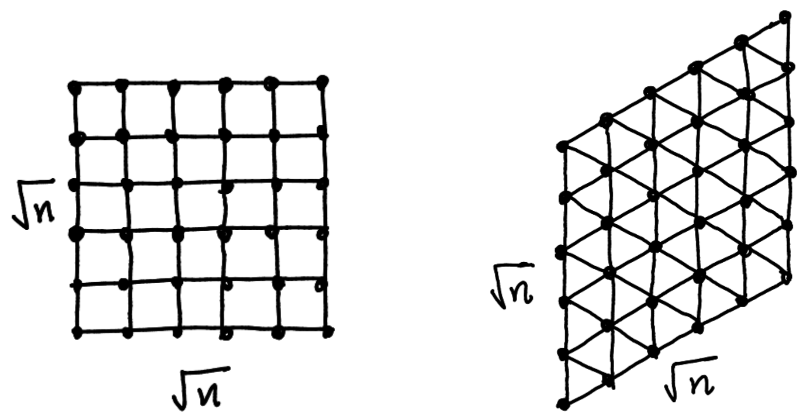

Figure 1: Unit distances in the square and triangular grids.

These constructions all give linear lower bounds. Can we do better? Erdős had a key idea: if we find a set of n points that define very few distinct distances, then by the pigeonhole principle, one of those distances must occur between many pairs. By rescaling the point set, we may assume this distance is a unit distance, giving a construction with many unit distances.

This gives rise to the following related problem. For a set $P \subseteq \mathbb{R}^2$ of points in the plane, let $d(P)$ denote the number of distinct distances between points in P, and let $d(n) = \min d(P)$, where the minimum is taken over all sets of n points in the plane. Our above reasoning with the pigeonhole principle then gives:

Observation 2.1.

$$u(n) \geq \frac{\binom{n}{2}}{d(n)}.$$

The regular n-gon shows $d(n) \leq \lceil \frac{n-1}{2} \rceil$, again giving a linear bound. However, the following result of Erdős shows this is not the correct order of magnitude.

Theorem 2.2 (Erdős, 1946). *There is a constant $C > 0$ such that*

$$d(n) \leq \frac{Cn}{\sqrt{\log n}}.$$

By our observation, this implies $u(n) \geq cn\sqrt{\log n}$ for some other constant $c > 0$.

Proof. We again consider the $\sqrt{n} \times \sqrt{n}$ (square) grid. Each point is of the form (x_1, x_2) with $x_i \in [\sqrt{n}]$. The distance formula gives us that for $(x_1, x_2), (y_1, y_2)$ we have $(x_1 - y_1)^2 + (x_2 - y_2)^2 \in [0, 2n - 1]$.

This already shows the number of distances is at most $2n$. However, we also see that the distances that occur are the sums of two squares of integers. It is a number theoretic fact that there are at most $\frac{Cn}{\sqrt{\log n}}$ sums of two squares in $[2n - 1]$, giving the upper bound on $d(n)$. □

Having shown that linear was not the correct order of magnitude for either problem, Erdős then conjectured that it was not too far off — that while you may be able to improve by a poly-logarithmic factor, you could not do *polynomially* better.

Conjecture 2.3 (Erdős, 1946). *When considering n points in the plane, we have:*

1. $d(n) = n^{1-o(1)}$.
2. $u(n) = n^{1+o(1)}$.

Note that, by the Observation, the second conjecture implies the first (but not the other way around).

These two natural questions of Erdős are more than seventy years old and had a tremendous impact on the field of combinatorial geometry. They are responsible for the introduction of several new methods to the field.

So far we have only considered constructions. What can we say about the other direction? Let us start with a lower bound on $d(n)$. Let P be an arbitrary set of n points and fix two arbitrary points $p_1, p_2 \in P$. Let ℓ_i be the number of distinct distances a point p_i participates in. Let us draw around p_i ℓ_i circles of these radii. Then each of the remaining $n - 2$ points are on one of the circles for both of these family of circles. There are $\ell_1 \ell_2$ pairs of circles from the two families and each pair can give rise to at most two intersection points. Since each of the reamining $n - 2$ points are one of these intersections, we have $2\ell_1\ell_2 \geq n - 2$. That means $d(P) \geq \max\{\ell_1, \ell_2\} \geq \sqrt{(n-2)/2} = \Omega(n^{1/2})$.

This was the original lower bound of Erdős. In the following decades the exponent has been improved several times, from 1/2 to 2/3 to 3/4, to 4/5, to 6/7, and so on. Finally in 2010 Guth and Katz have shown that $d(n) \geq \frac{cn}{\log n}$ for some $c > 0$, thus proving Conjecture 2.3(i).

As we have mentioned, the unit distance problem is supposed to be more difficult, as it implies the distinct distance problem, and indeed, the results obtained so far suggest that it is. We proved an upper bound of the order $n^{3/2}$ in a homework in the Turán number section, by showing that the (infinite) unit distance graph is $K_{2,3}$-free. This bound was very difficult to improve, with only some logarithmic factors being shaved off, until the following bound of Spencer, Szemerédi and Trotter.

Theorem 2.4 (Spencer-Szemerédi-Trotter, 1984). $u(n) \leq 4n^{4/3} + \frac{1}{2}n$.

Here we will prove this theorem with a surprising application of graph drawings. So first we take a detour to introduce our tool.

2.2 Planarity and Crossings

We first recall the definition of a drawing of a graph from Discrete Math I.

Definition 2.5. *A* drawing *of a graph $G = (V, E)$ consists of an injective map $\varphi : V \to \mathbb{R}^2$ together with, for every edge $e = \{u, v\} \in E$, a curve $\gamma_e : [0,1] \to \mathbb{R}^2$ with with $\gamma_e(0) = \varphi(u)$ and $\gamma_e(1) = \varphi(v)$.*

Recall that a planar graph is one that can be drawn without any pairs of edges crossing. We formally define a crossing in the following manner.

Definition 2.6. *A* crossing *of edges $e \neq f$ is a point in $\mathbb{R}^2$ that is not a common endpoint of e and f, but lies in $\gamma_e \cap \gamma_f$.*

Definition 2.7. *A graph G is* planar *if it has a drawing in the plane without any crossings.*

Euler's Formula, stated below, relates the number of faces, edges and vertices in a planar map. As we have seen in Discrete Math I, this bounds the number of edges in a planar graph.

Theorem 2.8 (Euler, 1751; Descartes, 1639)**.** *A connected planar map with v vertices, e edges, and f faces must satisfy $v - e + f = 2$.*

Corollary 2.9. *If G is an n-vertex planar graph with $m \geq 3$ edges, then $m \leq 3n - 6$.*

Proof. Consider a planar drawing (without crossings). If G is disconnected, add some edges to get a connected planar map. By Euler's Formula, $n - m + f = 2$.

Now we double count (edge, face) pairs where the edge is on the boundary of the face. Each edge is on the boundary of at most two faces, while each face has at least three edges on its boundary. This gives us that the number of (edge, face) pairs is at least $3f$ and at most $2m$, and so $f \leq \frac{2}{3}m$.

Substituting this into Euler's Formula, we have $n - m + \frac{2}{3}m \geq 2$, or $m \leq 3n - 6$. □

Thus a graph with at least $3n - 5$ edges contains crossings. We will be interested in the number of such crossings that must appear.

Definition 2.10. *The* crossing number *of a plane drawing of a graph is the sum over all pairs of distinct edges of the number of crossings between those pairs. The crossing number of a graph G, $cr(G)$, is the minimum crossing number of a drawing of G.*

Remark 2.11. *$cr(G) = 0$ if and only if G is planar.*

Note that we are counting crossings with multiplicity. If a point appears on many edges, we count one crossing for each pair of those edges. If a pair of edges crosses 100 times, they contribute 100 crossings.

For example, consider the cycle C_8, with all four diameters added. If these diameters are drawn within the circle, all crossing at the center of the cycle, that drawing would have crossing number $\binom{4}{2} = 6$. We could instead draw two diameters outside, and two inside, resulting in a crossing number of only two.

2.3 Lower bounds on the crossing number

Our main tool will be a strong lower bound on the crossing number of a graph. We begin with a first, weak, lower bound.

Claim 2.12. *If G is an n-vertex m-edge graph, $cr(G) \geq m - 3n$.*

Proof. Induction on m.

Base case: $m \leq 3n$ gives that $m - 3n \leq 0$, and the result follows trivially.

Induction step: $m > 3n$. Take an optimal drawing of G. By Corollary 2.9, this cannot be planar, and so there exist edges e, f that cross. Let $G' = G \setminus \{e\}$. By the inductive hypothesis, we have that $\mathrm{cr}(G') \geq m - 1 - 3n$. So, in our optimal drawing, there are at least $m - 1 - 3n$ crossings not involving e, together with at least 1 crossing involving e, and so $\mathrm{cr}(G) \geq m - 3n$. □

We now make some observations about where the crossings can take place in an optimal drawing of a graph.

Claim 2.13. *In an optimal drawing of G, there are no crossings between pairs of edges that share a vertex.*

Proof. Suppose not. Then there exists an optimal drawing where there exist edges $e = \{u, v\}$ and $f = \{u, w\}$ that cross. Let x be the last crossing of e and f, starting from u and heading towards v and w respectively. Let $\gamma_{e,1}$ be the curve of e from u to x, $\gamma_{e,2}$ be the curve of e from x to v,$\gamma_{f,1}$ be the curve of f from u to x, and $\gamma_{f,2}$ be the curve of f from x to w.

We switch the paths to obtain a better drawing of G. Let $\gamma'_e = \gamma_{f,1} + \gamma_{e,2}$, and $\gamma'_f = \gamma_{e,1} + \gamma_{f,2}$. Now, we slightly pull apart γ'_e and γ'_f at x. These curves no longer cross at x, implying the number of crossings decreases. This contradicts the optimal of the original drawing. □

Remark 2.14. *This proof also shows that in an optimal drawing, any two edges e, f cross at most once.*

With these preliminaries in place, we will prove a much stronger lower bound on the crossing number. The proof is a beautiful and surprising application of the probabilistic method, amplifying the bound of Claim 2.12.

Lemma 2.15 (Crossing number inequality)**.** *Let G be an n-vertex, m-edge graph, where $m \geq 4n$. Then $cr(G) \geq \frac{m^3}{64n^2}$.*

Remark 2.16. *This shows $cr(K_n) = \Theta\left(n^4\right)$, since we always have the upper bound $cr(G) = O\left(\binom{m}{2}\right)$. Determining the asymptotics of $cr(K_n)$ remains an open problem.*

Proof of Lemma. Fix an optimal drawing of G. Fix $p \in [0,1]$, whose value we shall determine later. Let G_p be the random induced subgraph of G, where each vertex survives independently with probability p.

Let X be the number of vertices in G_p, let Y be the number of edges in G_p, and let $Z = \mathrm{cr}(G_p)$. By Claim 2.12, $Z \geq Y - 3X$, and so $Z - Y + 3X \geq 0$. Thus, we have that $\mathbb{E}[Z - Y + 3X] \geq 0$. By the linearity of expectation, $\mathbb{E}[Z] - \mathbb{E}[Y] + 3\mathbb{E}[X] \geq 0$. Now, we calculate:

- $\mathbb{E}[X]$: np, as each of the n vertices remains with probability p.
- $\mathbb{E}[Y]$: The probability that any given edge survives is p^2. By linearity of expectation, we have that $\mathbb{E}[Y] = mp^2$. (Note that linearity of expectation is important here, since edges that intersect are not independent of one another.)
- $\mathbb{E}[Z]$: We can upper bound $\mathbb{E}[Z]$ by looking at the expected number of crossings that survive from our fixed optimal drawing of G. By Claim 2.13, the two edges of every crossing involve four distinct vertices, implying a crossing survives with probability p^4. By linearity of expectation, we have that $\mathbb{E}[Z] \leq \mathrm{cr}(G)\, p^4$.

Putting this all together,

$$0 \leq \mathbb{E}[Z] - \mathbb{E}[Y] + 3\mathbb{E}[X] \leq \mathrm{cr}(G)\, p^4 - mp^2 + 3np,$$

and so

$$\mathrm{cr}(G) \geq \frac{m}{p^2} - \frac{3n}{p^3}.$$

We now need to choose p to make this lower bound as large as possible. Some elementary calculus suggests setting $p = \frac{4n}{m}$. Note that this is at most 1 (and hence a valid probability), since $m \geq 4n$. Plugging this in gives

$$\mathrm{cr}(G) \geq \frac{m}{\left(\frac{16n^2}{m^2}\right)} - \frac{3n}{\left(\frac{64n^3}{m^3}\right)} = \frac{m^3}{64n^2}.$$

□

2.4 Back to unit distances

For our grand finale, we use the crossing number inequality to provide an upper bound on the number of unit distances in a set of n points in the plane.

Proof of Theorem 2.4. Fix an arrangement of n points with $u(n)$ unit distances. Observe that two points x_1 and x_2 are at unit distance if and only if x_2 belongs to the unit circle with center x_1, and vice versa. For each x_i, let d_i be the number of points on the unit circle around x_i, i.e., the number of unit distances it defines. Now, we can write

$$u(n) = \frac{1}{2}\sum_{i=1}^{n} d_i.$$

We can now define a graph G, whose vertices are the n points. For each x_i, we go around the unit circle around x_i, adding an edge in G between the neighboring pairs of vertices on the circle. For instance, if the vertices x_1, x_10, x_5 are at unit distance from x_3, appearing in that order on the unit circle centered at x_3, we add the edges $\{x_1, x_10\}$, $\{x_10, x_5\}$ and $\{x_5, x_1\}$ to G.

If there is only one point on the circle, we do not add any edges. If there are only two points on a circle, we only add one edge between them. Thus a circle with d_i points contributes at least $d_i - 1$ edges to G. It is possible that an edge comes from two different circles, but not three or more, since there cannot be three points that are all at unit distance from both endpoints of the edge. We remove any multiple edges from G, leaving us with

$$m := e(G) \geq \frac{1}{2}\sum_{i=1}^{n}(d_i - 1).$$

Thus, we have $m \geq u(n) - \frac{1}{2}n$, or $u(n) \leq m + \frac{1}{2}n$. If $m < 4n$, then $u(n) < 4n + \frac{1}{2}n \leq 4n^{4/3} + \frac{1}{2}n$, and we are done.

So we may assume $m \geq 4n$. By the crossing number inequality, we have $\mathrm{cr}(G) \geq \frac{m^3}{64n^2}$. However, we also have a drawing of G, using the circular arcs from the unit circles around the points x_i. In this drawing, every crossing is the intersection of two of these unit circles. There are n unit circles, each pair of which can contribute at most two crossings. Thus, $\mathrm{cr}(G) \leq 2\binom{n}{2} \leq n^2$. Hence $m^3/(64n^2) \leq n^2$, or $m \leq 4n^{4/3}$, which implies $u(n) \leq 4n^{4/3} + \frac{1}{2}n$, as required. $\square$