Math 503 (Discrete Geometry) Notes

Kenneth Moore
University of British Columbia

January → April, 2022

Lecture 1

Recall some basic identities,

1. $\binom{n}{k} = \frac{n!}{(n-k)!k!}$
2. $(x+y)^n = \sum_{i=0}^{n} \binom{n}{i} x^i y^{n-i}$
3. $\binom{n}{k} = \binom{n}{n-k}$
4. $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$

Now, using this we see

$$\sum_{i=0}^{n} (-1)^i \binom{n}{i} = (1-1)^n = 0$$

We can also get

$$\sum_{k=0}^{n} \binom{n}{k}^2 = \sum_{k=0}^{n} \binom{n}{k}\binom{n}{n-k} = \binom{2n}{k}$$

Now let's have fun.

Theorem 1.1 (Fermat)**.** *If p is a prime and a is an integer not divisible by p, then p divides $a^p - a$.*

Proof 1, induction. The base case of $a = 0$ is trivial. Suppose $a > 0$ and write $a = b + 1$. We have

$$\begin{aligned} a^p - a &= (b+1)^p - (b+1) \\ &= b^p + \binom{p}{1} b^{p-1} + \ldots + \binom{p}{p-1} b + 1 - (b+1) \\ &= b^p - b + \binom{p}{1} b^{p-1} + \ldots + \binom{p}{p-1} b \end{aligned} \tag{1.1}$$

which is divisible by p by the induction hypothesis. □

Proof 2, combinatorial. Consider necklaces with p beads and a colors. We create equivalence classes of necklaces so that two necklaces are in the same equivalence class if there is a rotation moving one to the other. Observe that the size of these equivalence classes is p. Thus there are a^p partition classes, and we subtract a for the monochromatic necklaces. □

A harder problem. Consider 'open necklaces' which is just a string of beads. We have k colors, and suppose each color paints an even number of beads. We have $2n$ beads. What is the fewest number of cuts we need to make to make two collections of pieces which each have an equal number of each coloured bead? For example, in the following necklace, we only needed two cuts, and then the middle part vs. the two end parts are two collections with the same number of colors. There

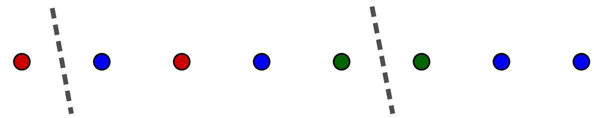

is a tricky proof which shows the answer is k.

Theorem 1.2 (Borsuk-Ulam)**.** *Let $f : \mathbb{S}^k \mapsto \mathbb{R}^k$ be a continuous map for an integer $k \geq 1$, and suppose $f(\overline{x}) = -f(-\overline{x})$. Then there exists $\overline{x} \in \mathbb{S}^k$ so that $f(\overline{x}) = 0$.*

To relate this to the necklace problem, we sort of just make a continuous version of the problem. Consider a line segment which is painted with different colors and each color set is measurable. In short, it turns out that the point that exists by the above theorem is a proper way to make cuts for each color. The idea is to find a color partition in each coordinate, and then find k cuts using the theorem. We define the functions

$$f_j = \sum_{i=1}^{k+1} \mathrm{sign}(x_i) m_j(i) \tag{1.2}$$

where $m_j(i)$ is the measure of the j^{th} color in the i^{th} coordinate[1].

Lecture 2

Book Recommendation: Enumerative combinatorics is a good book to check out. Polya-Szego problems are good as well. There is another book which is in the course syllabus by Stelle that you should check out. Richard Stanley is a book for just enumerative combinatorics, specialized. Winogradov had a good book of problems in number theory.

Notation,

1. $!n$ is 'derangements', the # of permutations without fixed points. For example,

$$\begin{matrix} 1 & 2 & 3 & 4 \\ 4 & 3 & 2 & 1 \end{matrix} \tag{2.1}$$

is a permutation with no fixed points. We can figure out a recursive statement which defines these numbers. If we choose 1 as a special element, there are two cases,

$$\begin{matrix} 1 & ... & i & ... \\ i & ... & 1 & ... \end{matrix} \tag{2.2}$$

In this case, we have $(n-1)\,(!(n-2))$. and

$$\begin{matrix} 1 & ... & i & ... & j & ... \\ i & ... & j & ... & ? & ... \end{matrix} \tag{2.3}$$

[1]See 3Blue1Brown https://youtu.be/yuVqxCSsE7c

In this case consider what happens if you just remove 1, and put i wherever 1 was on the bottom row. You will not get any fixed points this way, so we get $(n-1)(!(n-1))$ more permutations for a total of $(n-1)\,(!(n-2)+!(n-1))$.

Recall that every permutation can be represented by the cycles. Like in (2.1), the cycle was (1423). We define the Stirling numbers of the first kind, $s(n,k) = \#$ of permutations of n elements with k cycles. Can we create a recursion for these?

$$s(n,k) = \begin{bmatrix} n+1 \\ k \end{bmatrix} = n \begin{bmatrix} n \\ k \end{bmatrix} + \begin{bmatrix} n \\ k-1 \end{bmatrix} \tag{2.4}$$

If 1 is a fixed point, you are left with $n-1$ elements, and 1 is a cycle, so this case results in $s(n,k-1)$, which is where we get the second term in (2.4). The first term comes from a similar argument to before, removing 1 and then adding it back anywhere.

Define Stirling numbers of the second kind $S(n,k) = \left\{ {n \atop k} \right\}$ are the number of partitions of an n-element set into k non-empty classes. Note that two partitions into classes are distinct if there are two elements which are in the same class in one partition and in different classes in the second partition. Writing this as a recursion, we have

$$\left\{ {n+1 \atop k} \right\} = k \left\{ {n \atop k} \right\} + \left\{ {n \atop k-1} \right\} \tag{2.5}$$

Where the second one comes from partitioning the new element into its own set.

2.1 Graphs

What is the maximum number of edges in a graph with v vertices if there are no crossings? It turns out that there at most $3v-6$ for planar graphs. The follows from the fact that the Euler characteristic is 2

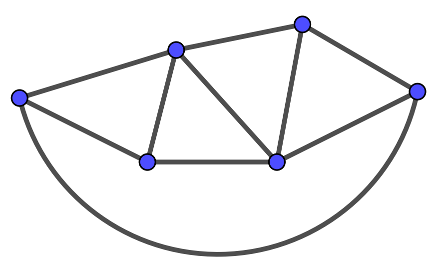

This graph for example has $v = 6$, $e = 10$ and $f = 6$. By Euler's formula, we have $v-e+f = 2$ and luckily our graph satisfies it! We talked a lot about colourings here, nothing rigorous though.

If G is a connected planar graph, then in its embedding to the plane, $v - e + f = 2$. This can be proven with induction on e. This formula leads to the bound on the number of edges. We have $f = f_3 + f_4 + ... + f_k$ where f_j is the number of faces with j edges adjacent. So consider $3f_3 + 4f_4 + ... + kf_k = 2e$ since we count each edge twice, once from each side. We then have $3f \leq 2e$, so then $f \leq \frac{2}{3}e$. By Euler's formula, we have $v - e + \frac{2}{3}e \geq 2$, and then $e \leq 3v-6$.

Lecture 3

Here we did a little graph theory. We generally say $V(G_n)$ is the vertex set and $E(G_n)$ is the edge set of the graph G_n, and $n = |V|$, $e = |E|$.

Definition 3.1. *A connected graph on n vertices with $n-1$ edges is called a tree, denoted T_n. Equivalently, a connected graph with no cycles.*

Handshake Lemma is useful to see these are equivalent, $\sum_{v_i \in V(G_n)} \deg(v_i) = 2e$.

Theorem 3.2 (Cayley's formula)**.** *The number of labelled trees on n vertices is n^{n-2}.*

Proof. We show this by Prüfer codes. That is, we show every sequence of $n-2$ numbers from n digits corresponds to a unique tree. For example,

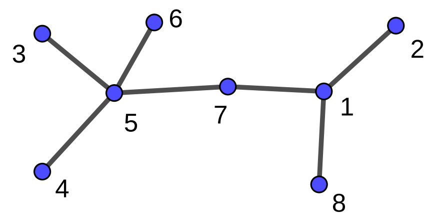

To do this,

- Find the smallest labelled leaf. In this case, it's 2. Write down its neighbours label in the sequence.
- Repeat this until there are just two vertices left.

Now, this is a well defined process, we do not make any decisions about what do do at any step. The above tree for example gives you $(1,5,5,5,7,1)$

Going backwards, we can recover the tree like this. Write the leaves (the numbers that don't appear in the sequence) under the sequence.

$$\begin{array}{cccccc} 1 & 5 & 5 & 5 & 7 & 1 \\ 2 & 3 & 4 & 6 & 8 & \end{array} \tag{3.1}$$

And eliminate numbers one by one, connecting top to the smallest number on the bottom and adding a number to the bottom when it becomes a leaf (no longer appears on top),

$$\begin{array}{cccccc} \cancel{1} & \cancel{5} & \cancel{5} & \cancel{5} & 7 & 1 \\ \cancel{2} & \cancel{3} & \cancel{4} & \cancel{6} & 8 & 5 \end{array} \tag{3.2}$$

then in the last step, connect the two leaves that are left on the bottom, in this case 1 and 8.

$$\begin{array}{cccccccc} \cancel{1} & \cancel{5} & \cancel{5} & \cancel{5} & \cancel{7} & \cancel{1} & & \\ \cancel{2} & \cancel{3} & \cancel{4} & \cancel{6} & 8 & \cancel{5} & \cancel{7} & 1 \end{array} \tag{3.3}$$

How do we know a Prüfer code represents a tree? This is because we can get the degree sequence back. $\square$

We call C_n a cycle of length n. Also, P_n the path of length $n-1$, which has $n-1$ edges. Also, K_n is the complete graph on n vertices which has $n(n-2)/2$ edges. Also, $K_{n,m}$ is the complete bipartite graph, which has two classes V_1 of n and m vertices, and the graph has all the edges connecting vertices between classes and none within classes.

The chromatic number of G_n is the least number so that the vertex set of G_n can be coloured so that no edge connects the same color.

Theorem 3.3 (Ramsey's)**.** *For any k there is $n_0(k)$ such that if $n \geq n_0(k)$ then for any two-colouring of the edges there will be a mono-chromatic complete subgraph $\widetilde{K}_k$.*

Note that $R(3) = n_0(3)$ is 6. You can show with a nice argument that this is enough, and prove Ramsey's theorem in a similar way, inductively. In fact $R(k) \leq 4^k$. In fact this is the best known bound! (WAS* AT THE TIME!!)

Lecture 4

4.1 Matchings in Graphs

Definition 4.1. *A matching is a subgraph so that no edge shares a vertex. A perfect matching is a matching with $\frac{n}{2}$ edges in G_n.*

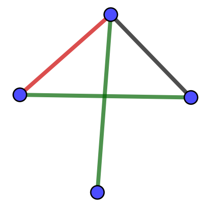

A common question is how many perfect matchings does a certain graph have? Here are two different ones (red and green) for the Petersen graph.

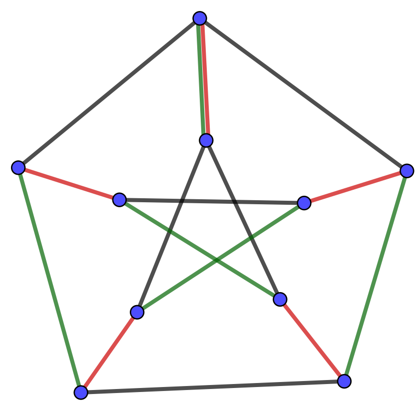

Exercise 4.2. *How many perfect matchings in L_{2n}? (This is a graph with n boxes.)*

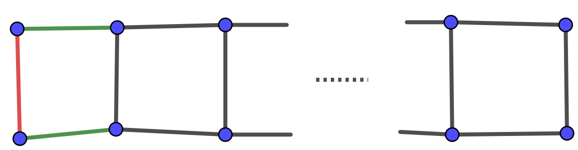

Proof. The key to this is Fibonacci. If we start with the vertical edge (red), then we basically remove the first box, so we get $PL(n-1)$ in this case. If we start with the top and bottom edges in the first box (green), we remove the last two boxes, so we get $PL(n-2)$ in this case. So in total, $PL(n) = PL(n-1) + PL(n-2)$. We then show the initial values are the same as the Fibonacci, so the values are the same. □

Recall the value $R(k)$. We saw last time that any two colouring of the edges of $K_{R(k)}$ results in a monochromatic subgraph (K_k). Let $R(\overbrace{3, 3, \dots, 3}^{k})$ be the number such that the k colouring of the edges of $K_{R(3,3,\dots,3)}$ results in a monochromatic triangle. we can prove the result for this similar to what we did last time, choose any vertex and take the most popular color among edges, say, red in this case. The size of edges with the same color is $\geq \frac{n-1}{k}$, and suppose there is no triangle; then there is no more of that color in the subgraph. We then have $R_k(3) \leq kR_{k-1}(3)$, and so $R_k(3) \lesssim k!$.

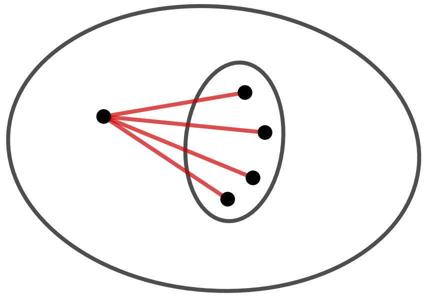

We can use this to prove Schur's theorem; for any k there is a bound $S(k) \in \mathbb{N}$ such that for any k-colouring of $[S(k)]$, there is a monochromatic solution to $x + y = z$. Schur was interested in this question because if it was false, he could prove Fermat's last theorem. What this theorem implies is that $x^n + y^n = z^n \mod p$ has a solution for any n, as long as p is a large enough prime. In $\mathbb{F}_p$, $A_1 = \{a^n \mid a \in \mathbb{F}_p\}$, $A_b = \{ba^n \mid a \in \mathbb{F}_p\}$.

Another technique, consider a graph with $n+1$ vertices, labelled from 1 up to $n+1$. Make it a complete graph, and color the edge connecting $j \geq i$ with the color of the vertex $j - i$. If there is a triangle in that graph, then $j - i$ and $\ell - j$ and $\ell - i$ all have the same color. Letting those values be x, y, and z, we have $x + y = z$ all having the same color.

Conjecture 4.3. *for every k, there is an integer $F(k) := n$ such that there are numbers in $[n]$ coloured with k colors, such that x, y, $x + y$, xy all have the same color.*

4.2 Normal Generating Functions.

First example, consider $a_j = 1$ for all $j \in \mathbb{N}$. The generating function GF is a formal power series $GF(x) = \sum_{i=0}^{\infty} a_i x^i$. In this fist case, this is simply $\frac{1}{1-x}$. Note we often ignore radii of convergence.

Next example, let $a_n = a_{n-1} + 2a_{n-2}$, $n \geq 2$ and $a_0 = a_1 = 1$. As $n \to \infty$, we should see $\frac{a_n}{a_{n-1}} \to$ some number α, so

$$\alpha \approx \frac{a_n}{a_{n-1}} = 1 + 2\frac{a_{n-2}}{a_{n-1}} = 1 + \frac{2}{\alpha} \tag{4.1}$$

so $\alpha^2 - \alpha - 2 = 0$. The roots are 2 and 1. Then $a_n = A2^n + B(-1)^n$. Using the initial values, $A = \frac{2}{3}$ and $B = \frac{1}{3}$. Thus $a_n = \frac{1}{3}(2^{n+1} + (-1)^n)$. This method was a little less scientific than we would like, mostly just because we made an assumption of convergence. Let's use generating functions. We have

$$S(x) = \sum_{n=0}^{\infty} a_n x^n \tag{4.2}$$

and try to find $S(x)$. To do this, try multiplying by $-x$ and by $-2x^2$, and we have

$$\begin{aligned} S(x) &= a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \dots \\ -xS(x) &= -a_0 x - a_1 x^2 - a_2 x^3 + \dots \\ -2x^2 &= -2a_0 x^2 - 2a_1 x^3 + \dots \end{aligned} \tag{4.3}$$

and check it out, the x^2 terms in these three sum to 0. Same with every term higher than x. So all we have is $(1-x-2x^2)S(x) = 1$, and so $S(x) = \frac{1}{1-x-2x^2}$. After partial fractions, we have $S(x) = -\frac{1}{3}\left(\frac{1}{2x-1} - \frac{1}{x+1}\right)$

Next example, consider $a_n = a_{n-1} + 2^n$. To do this, you first have to find the generating function of 2^n. It is $\frac{1}{1-2x}$.

4.3 Exponential Generating Functions.

The exponential generating function is

$$F(x) = \sum_{i=0}^{\infty} a_i \frac{x^n}{n!} \tag{4.4}$$

So once again, let's see the function for $a_i = 1$? It's of course just e^x. If it was alternating, $a_j = (-1)^j$, we have e^{-x}. Something a bit less simple, let

$$S_E(x) = \sum_{i \text{ is even}} \frac{x^i}{i!}, \quad 1\ 0\ 1\ 0\ 1\ 0\ ...$$

or

$$S_O(x) = \sum_{i \text{ is odd}} \frac{x^i}{i!}, \quad 0\ 1\ 0\ 1\ 0\ 1\ ...$$

One thing is for sure, if we have $S_E + S_O$, we have e^x. We also see that $S'_E(x) = S_0(x)$ just from the sums. We thus have $S_E + S'_E = e^x$, we could also get $S_E(x) - S''_E(x) = 0$. Solving this differential equation is not hard, and we get $\frac{1}{2}\left(e^x + e^{-x}\right) = \cosh(x)$.

Lecture 5

5.1 Graphs and their Substructures

Take a graph G_n and suppose $e \geq cn^2$. We can make a graph with about $n^2/4$ edges that has no triangle. This is the complete bipartite graph.

Theorem 5.1 (Mantel's)**.** *In a graph containing no triangle, the maximum number of edges is* $\frac{n^2}{4}$.

More generally, $ex(G,n)$ is the max number of edges without G as a subgraph. For example, $ex(C_4,n)$? It turns out that $ex(C_n,n) \leq cn^{3/2}$. We convert this into a problem of counting cherries (P'_3s), since two vertices cannot share two neighbouring vertices otherwise you have a triangle. The question is now, how many cherries are in a G_n with e edges? If degree of a vertex is d, we have $\binom{d}{2}$ cherries. Thus

$$\#cherry = \sum_{v \in V(G_n)} \binom{\deg(v)}{2} \tag{5.1}$$

We also know that $e = \frac{1}{2}\sum_{v \in V(G_n)} \deg(v)$. These are clearly related, but it might be hard to see how. Let's use Cauchy-Schwartz on this. More accurately, we use Jensen's inequality.

$$\frac{1}{2}\sum_{v \in V(G_n)} \deg(v)^2 \approx \#cherry, \qquad \sum_{v \in V(G_n)} \deg(v) = 2e \tag{5.2}$$

The average degree is $\frac{2e}{n}$, so

$$\sum \deg^2 \geq n\left(\frac{2e}{n}\right)^2 \approx \frac{e^2}{n} \tag{5.3}$$

So we have about $\frac{e^2}{n}$ cherries. If there is no C_4, then the number of cherries is less than # of pairs of vertices. Then $\frac{e^2}{n} \lesssim n^2$ and so $e \leq cn^{3/2}$.

For the lower bound on $ex(C_4.n)$, we do a construction and see that it is also $cn^{3/2}$. Let v_1 be the set of points, and v_2 the set of lines. Make a graph on these elements, connecting a point to a line if the line passes through the point. This turns it into an incidence geometry problem. Now, what about $ex(n, P_4)$? $ex(K_{s,t}, n)$ is an open problem and an important one. Note that for any G_n with e edges, there is a bipartite subgraph $G(A, B)$, $|A| \approx |B| \approx \frac{n}{2}$ with $\approx \frac{e}{2}$ edges.

A proof strategy for $ex(K_{s,t}, n)$. You can find large sets of numbers so that the sum of a number from one and the other always gives you a prime.

First, count the number of multi-cherries with s berry. That is, $\sum_{v \in V(G_n)} \binom{\deg(v)}{s}$. If there are no $K_{s,t}$ in G_n, then the number of s-multicherries is $\leq \binom{n}{s}(t-1)$ since there are not t vertices sitting on any one multicherry. Similar to before we get $n\left(\frac{e}{n}\right)^s \leq n^s t$, and then $e \leq cn^{2-\frac{1}{s}}t^{\frac{1}{s}}$ which is also the conjectured lower bound (Kővári-Sós-Turán).

Book recommendation: The Cauchy-Schwartz Masterclass.

Lecture 6

Definition 6.1. *A graph G_n is* dense *if the number of edges is $c \cdot n^2$. G_n is* sparse *if $e = o(n^2)$.*

We quickly summarized the material from last time, notably that if e is large enough in G_n, then it contains $K_{s,t}$. We looked at two theorems about sparseness,

Theorem 6.2. *Every* dense *graph on n vertices has a complete bipartite graph subgraph $K_{\log n, \log n}$.*

Theorem 6.3. *If G_n contains e edges, then there is $G'_m \subseteq G_n$ where every vertex has degree at least half the average degree in G_n*

Proof. Choose any $v_i \in G_n$ with degree $deg(v_i) < \frac{1}{2}avg = \frac{e}{n}$ and remove it. It's not possible to remove all the vertices, because each time we remove less than the average number of edges, so if we removed all the vertices, we would be left with some floating edges!

So there is a substructure where all the remaining vertices have $deg > \frac{1}{2}avg$ and the number of edges removed is $\leq n(\frac{1}{2}avg - \varepsilon) < e$. If for example G_n is dense, it has cn^2 edges and thus it has a subgraph where every vertex has degree $\geq cn$. $\square$

We now talked a bit about adjacency matrices. For a graph on 2 sets $G(A, B)$, we put one set of vertices on one axis and the other on the other. If a graph has a bipartite subgraph $K_{s,t}$, it means that there is a submatrix of size s, t that is all ones.

Consider $|A| = |B| = n$, and the min degree $\geq c_n$, and remove rows of the adjacency matrix based on the maximum number of ones in a row. Each time we remove a row, we reduce the degree by a positive fraction of c_n. In conclusion, every dense graph contains $K_{t,t}$ with $t \approx \log n$. If G_n has $n^{2-\varepsilon}$ edges and $\varepsilon \to 0$ as $n \to \infty$, then G_n contains $K_{t,t}$ where $t \to \infty$.

A Note: I was thinking about the homework due tonight; may have missed some things.

Lecture 7

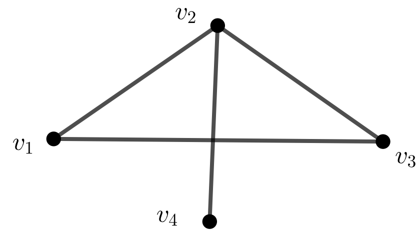

We defined the adjacency matrix. Recall that it is symmetric, so all its eigenvalues are real. For example, the graph above has adjacency matrix

$$\begin{matrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{matrix} \tag{7.1}$$

We usually use J to represent a matrix of all ones, and e the vector of all ones. e_j on the other hand is still just the j^{th} coordinate vector. Note that $A(G_n) \cdot e$ gives you the degree vector.

A useful application of these matrices is counting walks, and there are many extensions to cases of weighted and directed graphs. How many walks are there in G_n above, starting at v_3? A walk is a sequence of vertices where consecutive vertices are connected by an edge. So v_3, v_2, v_1 is a walk for example, but v_3, v_2, v_3 is too. Thus there are 5 in this graph. We define $w_k(i, j)$ to be the number of walks from vertex i to j of length k. A question is, what is $w_4(3, 1)$ in the above graph? The adjacency matrix helps. In A^k, the coefficient $a_{ij} = w_k(i, j)$.

Another question, what are $\text{tr}(A^2)$, $\text{tr}(A^3)$, and $\text{tr}(A^4)$? The diagonal of A^2 just counts the degree of each vertex. $\text{tr}(A^3)$ is the number of triangles in the graph. $\text{tr}(A^4)$ is more interesting...

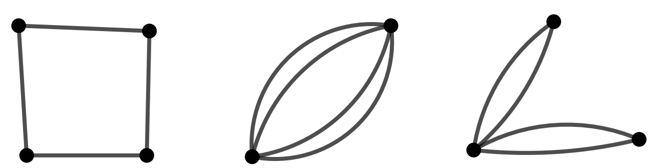

It counts the number of all these things summed. The number of C_4s this counts is not just all C_4s in the graph, but all homomorphisms of them too, because it counts starting at any of the 4 vertices in a C_4 and either direction you can start out in. We can check if the graph is connected in this way too, let $B = A + A^2 + ... + A^{n-1}$. Then G_n is connected if and only if B has positive entries only.

The spectrum of a graph is $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$, the eigenvalues of the adjacency matrix. There are eigenvectors of each of these values, and if x and y are eigenvectors corresponding to distinct eigenvalues, then they are orthogonal since A is a nice symmetric matrix. Thus the vectors form an orthonormal basis of size up to n.

A graph G_n is said to be d-regular if $\deg(v_i) = d$ for all $q \leq i \leq n$. For example, the Petersen graph has degree 3 on all vertices. What do we know about the spectrum of d-regular graphs? One thing is that $Ae = de$, so d is an eigenvalue. A claim is that all other eigenvalues $|\lambda| \leq d$. To see this, let v be any other eigenvector, and find its largest entry. we can scale the vector to have largest entry of 1. We can see the largest entry in $A \cdot v$ will have largest entry strictly less than d, and so $|\lambda| < d$.

Theorem 7.1 (Geřshgorin theorems)**.**

- *In a complex matrix M, if for all i, $|a_i i| > \sum_{j \neq i} |a_i j|$, then M has full rank.*
- *Let $S = \sum_{i \neq j} |a_{ij}| = r_i$. Every eigenvalue is in one of the Geřshgorin disks,*

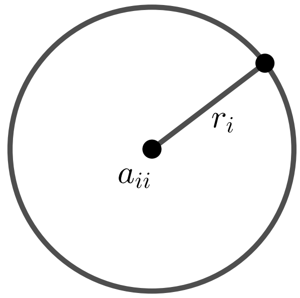

- *If $\Delta(G_n) = \Delta$ is the max degree, then $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$. Then $d_{avg} \leq |\lambda_n| \leq \Delta$.*

A fact we will apparently use a lot is that since A is a symmetric matrix, it has an orthonormal basis of eigenvectors.

Suppose we have two subsets of vertices U and W. How can we use the adjacency matrix to count the number of edges between U and W? To do this we use indicator vectors of each set. Multiply the matrix between the two.

Lecture 8

There is another kind of graph matrix, called the Laplace matrix $D - A$ where D is the diagonal matrix with $d_{ii} = \deg(v_i)$. There is also the unsigned Laplace matrix which is $D + A$. The benefit of these is that they are positive definite matrices!

Recall that A^k gives you the number of walks from v_i to v_j. Let G_n be a connected simple graph with diameter (the largest minimum distance between two points in the graph) d. Claim: G_n has at least $d + 1$ distinct eigenvalues.

Suppose there are t distinct eigenvalues. Then write $(A - \lambda_1 I)(A - \lambda_2 I)...(A - \lambda_t I) = 0$. Thus A^t is a linear combination of A^{t-1}, $A^{t-2}, \ldots, I$. Note that if $t \leq d$, then there is a 0 in position (i, j) of A^t if $dist(v_i, v_j) = d$.

8.1 Finding eigenvalues: Circulant matrices

C is a circulant (complex) matrix if it has the form

$$\begin{matrix} c_0 & c_{n-1} & c_{n-2} & .. & c_1 \\ c_1 & c_0 & c_{n-1} & .. & c_2 \\ c_2 & c_3 & .. & .. & .. \\ .. & .. & .. & .. & .. \\ c_{n-1} & c_{n-2} & .. & .. & c_0 \end{matrix} \tag{8.1}$$

The eigenvectors here are the roots of unity. if $\omega_j = \exp\left(\frac{2\pi i j}{n}\right)$ for $0 \leq j \leq n-1$ is the j^{th} root of unity. The eigenvalues are then $\lambda_j = \sum_{i=0}^{n-1} c_i \omega_j^i$

8.2 Finding eigenvalues: Strongly regular graphs

In a d-regular graph, d is an eigenvalue with eigenvector e. The Petersen graph for example is 3-regular, and has $girth(p) = 5$ (the smallest cycle is 5). Fun fact, the number of spanning trees in the Petersen graph is exactly 2000??

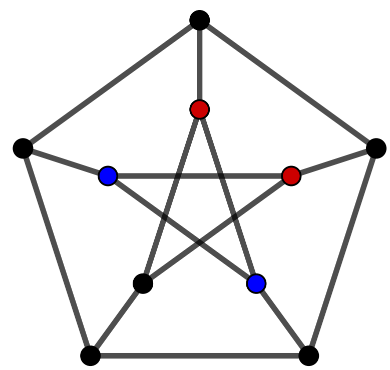

A graph is strongly regular if between any two vertices, the number of common neighbours depends only on whether the vertices are connected. For example, in the Petersen graph, any two that are not connected (like the red points) have exactly 1 common neighbour. Any two that are connected (like the blue) have 0.

Now, how can we find the eigenvalues of such a graph? A is the adjacency matrix of P. We know something about A^2. Taking $A^2 + A + I - D = \mathcal{J}$ (the all one matrix). In this case, we have $A^2 + A - 2I = \mathcal{J}$. Multiplying both sides by an eigenvector x which has eigenvalue λ then,

$$\lambda^2 x + \lambda x - 2x = 0 \tag{8.2}$$

this is because x is orthogonal to the all one vector e, since e is itself an eigenvector. We see λ is an eigenvalue of A^2 as well. Thus the eigenvalues satisfy $\lambda^2 + \lambda - 2 = 0$ and so $\lambda = 1, -2$. To recover the multiplicity, we use the fact that the trace of A is 0, so the sum of the eigenvalues is also 0. We can therefore solve for the multiplicities, in particular 1 has multiplicity 5, -2 has multiplicity 4, 3 has multiplicity 1.

Theorem 8.1. *There is no decomposition of the edges of K_{10} into three copies of the Petersen graph.*

Proof. Suppose $A + B + C = \mathcal{J} - I$ and A, B, C are adjacency matrices of Petersen graphs. Let V_A and V_B be the subspaces of the eigenvectors of the eigenvalue 1 in A and B. Both of these are 5 dimensional, and both orthogonal to e. The space $e^\perp$ is 9 dimensional, so there is some vector $x \in V_A \cap V_B$, so $\mathcal{J}x = 0$. Now, what happens if we multiply C by this? We get

$$Cx = (\mathcal{J} - I - A - B)x = -x - Ax - Bx = -3x$$

This means C has an eigenvector x with eigenvalue -3. Petersen graphs do not have that eigenvalue, so this is not possible. $\square$

Lecture 9

If a matrix is symmetric, then its eigenstuff is "nice"; all eigenvalues are real, and there is an orthonormal basis $v_1, v_2, \dots, v_n$ of the space of eigenvectors. $|v_i| = 0$, and $v_i \cdot v_j = 0$ for $i \neq j$.
Let A be the adjacency matrix of G_n. We can write

$$A = \sum_{i=1}^{n} \lambda_i v_i v_i^T \tag{9.1}$$

A symmetric minor B of A is a submatrix of A obtained by deleting some rows and the corresponding columns.

Theorem 9.1 (Interlacing eigenvalues)**.** *If A is an $n \times n$ symmetric matrix with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$ and B is an $n-k \times n-k$ symmetric minor of A with eigenvalues $\mu_1 \geq \dots \geq \mu_{n-k}$. Then $\lambda_i \geq \mu_i \geq \lambda_{i+k}$.*

In d-regular graphs, the largest eigenvalue $d = \lambda_1 \geq |\lambda_i|$. Also, λ is the second largest eigenvalue if $\lambda \geq |\lambda_i|$ for any $i > 1$. In a d-regular G_n, the second largest eigenvalue satisfies $\lambda \gtrsim \sqrt{d}$. The proof of this uses the facts that $\operatorname{tr} A^2 = dn$ and $\operatorname{tr} A^2 = \sum_{i=1}^{n} \lambda_i^2$.

Lemma 9.2. *The trace is the sum of the eigenvalues.*

Proof. The characteristic polynomial is

$$p(t) = \det(A - tI) = (-1)^n \left(t^n - (\operatorname{tr} A)t^{n-1} + \dots + (-1)^n \det A\right) \tag{9.2}$$

To see why that coefficient is $\operatorname{tr} A$, consider the cofactor expansion of $M = A - tI$ along any row i, any term involving an off-diagonal element M_{ij} eliminates both $a_{ii} - t$ and $a_{jj} - t$ so any such term cannot have high enough degree, and so the coefficient of t^{n-1} must come from $(a_{11} - t)\dots(a_{nn} - t)$. There is another result where a polynomial $x^n + a_1 x^{n-1} + \dots + a_k$ satisfies $a_1 = \sum$roots, whose application completes the proof. □

d-regular graphs with second largest eigenvalue close to $\sqrt{d}$ are called Ramanujan graphs.

Exercise 9.3. *If we have a k-regular graph with adjacency matrix A and two subgraphs of size s and t, the number of edges between S and T would be approximately $\sim \frac{k \cdot s \cdot t}{n}$. Estimate the error on this guess by*

$$\left| e(S,T) - \frac{k \cdot s \cdot t}{n} \right| \leq \lambda \sqrt{st} \tag{9.3}$$

which is a bound on the 'edge discrepancy'. This is also called a Cheeger-type inequality.

Proof. As usual, let $\lambda_1, \dots, \lambda_n$ be the n eigenvalues of the graph and let $v_1, \dots, v_n$ be an orthonormal basis of corresponding eigenvectors, with $\lambda_1 = k$ and $v_1 = \frac{1}{\sqrt{n}} e$. Use $\mathcal{J}$ to denote the matrix consisting of all 1s.

Let χ_S be the indicator vector for the set S (entry i is 1 if the i^{th} vertex is in S and 0 otherwise) and χ_T the indicator vector for T. We have $e(S,T) = \chi_S^\intercal A \chi_T$ and $\frac{k \cdot s \cdot t}{n} = \frac{k}{n} (\chi_S^\intercal \mathcal{J} \chi_T)$, thus

$$\left| e(S,T) - \frac{k \cdot s \cdot t}{n} \right| = \left| \chi_S^\intercal \left(A - \frac{k}{n} \mathcal{J} \right) \chi_T \right| \tag{9.4}$$

Now define $B = A - \frac{k}{n}\mathcal{J}$. Note that v_1 is also an eigenvector of $\mathcal{J}$ having the eigenvalue n. Since each v_i for $v_i \geq 1$ is perpendicular to e, they are also each eigenvectors of $\mathcal{J}$ with eigenvalue 0. Therefore the two matrices share all eigenvectors, so the eigenvalues of B are the differences of eigenvalues of the two, namely $0, \lambda_2, \dots, \lambda_n$. Now B is a real symmetric matrix, thus it is self-adjoint; we have

$$\|B\chi_T\|^2 = \langle B\chi_T, B\chi_T\rangle = \langle \chi_T, B^2\chi_T\rangle$$

Using the orthonormal basis v_i, we write the vector χ_T in coordinates to obtain

$$\langle \chi_T, B^2\chi_T\rangle = \langle \chi_T, \sum_{i=1}^{n} B^2\langle \chi_T, v_i\rangle v_i\rangle = \langle \chi_T, \sum_{i=2}^{n} \lambda_i^2\langle \chi_T, v_i\rangle v_i\rangle = \sum_{i=2}^{n} \lambda_i^2 \langle \chi_T, v_i\rangle^2 \leq \lambda^2 \|\chi_T\|^2$$

So in summary, $\|B\chi_T\| \leq \lambda \|\chi_T\|$. By the Cauchy-Schwarz inequality, we have

$$|\chi_S^\intercal B\chi_T| \leq \|\chi_S\| \, \|B\chi_T\| \leq \lambda \|\chi_S\| \, \|\chi_T\| = \lambda\sqrt{st} \tag{9.5}$$

and with equation (9.4), this completes the proof. □

Recall that $A = \sum \lambda_i u_i^\intercal u_i$, and let χ_S, χ_T be the indicator functions of S and T, so $\chi_S = \sum \alpha_i u_i$ and $\chi_T = \sum \beta_i u_i$. Then

$$e(S,T) = \chi_S^\intercal A \chi_T = \sum \alpha_i \lambda_i \beta_i$$

The first term in this sum ends up being $\frac{k \cdot s \cdot t}{n}$, and the rest of the terms in the sum are the error terms. We can also estimate them, using the second largest eigenvalue,

$$\left|\sum_{i>1} \alpha_i \lambda_i \beta_i\right| \leq \lambda \left|\sum_{i>1} \alpha_i \beta_i\right| \tag{9.6}$$

Lecture 10

Let A, B be sets. Denote by $A \times B$ the set of tuples (a, b),for $a \in A$ and $b \in B$. If $A, B \subset \mathbb{R}$, we can represent the elements of $A \times B$ as points in the plane $\mathbb{R}^2$. In this case $A \times B$ is a Cartesian product. Let us suppose $|A| = n$, $|B| = m$, finite sets of real numbers. The Cartesian product of them lies on a grid.

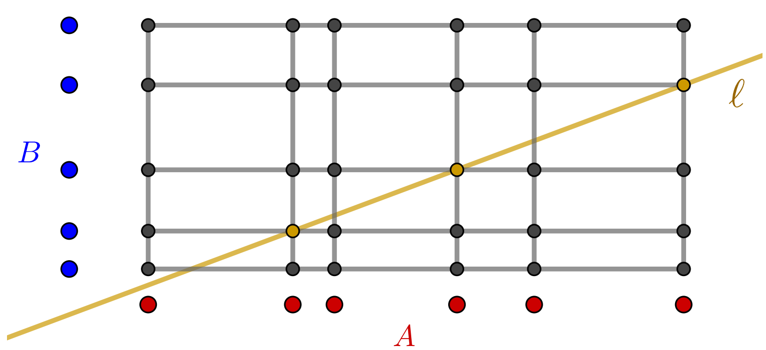

For now, suppose $|A| = |B| = n$. Let's give an upper bound on the k-rich lines on $A \times B$. A line ℓ is k-rich if $|\ell \cap (A \times B)| \geq k$. Let us denote the number of k-rich lines by x_k.

Partition $A \times B$ into "cells". We use $\frac{k}{10}$ vertical and $\frac{k}{10}$ horizontal separators so that there are $\sim \frac{10|A|}{k}$ points of the grid between each vertical separator, $\sim \frac{10|B|}{k}$ points of the grid between each horizontal separator. This results in cells which have about $\frac{100n^2}{k^2}$ points.

Now, given a line, two points in $\ell \cap (A \times B)$ that are consecutive are called a lucky pair if called lucky if they lie in the same cell. We have $\frac{k}{5}$ possible 'destroyers', that is, places to cut two consecutive points. Thus the number of lucky pairs on a k-rick line is $\sim \frac{4}{5}k$. This means we have at least $x_k \cdot \frac{4k}{5}$ lucky pairs are created with x_k k-rich lines. The number of lucky pairs across all lines is bounded above by

$$\frac{k^2}{100}\binom{\frac{100n^2}{k^2}}{2} \sim \frac{100n^4}{2k^2}$$

and therefore $x_k \leq c\frac{n^4}{k^3}$.

Theorem 10.1 (Szemerédi-Trotter)**.** *Given n points in $\mathbb{R}^2$, the number of unique lines contain k of the points is at most $O(\frac{n^2}{k^3})$.*

A famous problem in number theory is the sum-product problem. If A is a subset of $\mathbb{R}$ (or $\mathbb{C}$, $\mathbb{N}$, $\mathbb{F}_p$, etc.) then either $A + A$ or $A \cdot A$ should be large. In the real case, if $|A| = n$, $A \subset \mathbb{R}$, then we will show $|A + A| + |A \cdot A| \geq c \cdot n^{5/4}$ (although, the current best is an exponent slightly larger than $n^{4/3}$).

To do this, create the grid again with $A + A$ on one edge and $A \cdot A$ on the other. We now define lines by $\ell_{ij} = a_i(x - a_j)$. For every pair, we get a distinct line, so there are n^2 of these lines. Then these lines are all n-rich, as ℓ_{ij} contains $(a_k + a_j, a_i a_k)$ for any k. Using the previous theorem now, we see $n^2 \leq c\frac{(|A+A||A \cdot A|)^2}{n^3}$, so then $cn^{5/2} \leq |A + A|\,|A \cdot A|$.

In $\mathbb{C}$ or $\mathbb{F}_p$, there are some problems. In $\mathbb{C}^2$, lines don't cut the plane. In $\mathbb{F}_p$, there can be no partitioning help since there is no ordering. However, we will show something in this case. We consider $A \subset \mathbb{F}_p$ with $|A| << p$ and $\sqrt{p} \geq |A|$. Using graph spectra and a Cheeger-type bound, we can show that

$$|A + A|\,|A \cdot A| > c \min\left\{p\,|A|\,, \frac{|A|^4}{p}\right\} \tag{10.1}$$

and if $|A| \approx p^{2/3}$, then $\max\{|A + A|\,, |A \cdot A|\} > c\,|A|^{5/4}$. To do this we define the sum-product graph G_{SP}. We define $V(G_{SP}) = \mathbb{F}_p^* \times \mathbb{F}_p$. Two vertices $u = (a, b)$ and $v = (c, d)$ are connected by an edge iff $ac = b + d$.

Lecture 11

Let $\mathbb{F}_p$ be a finite field, with multiplicative subgroup $\mathbb{F}_p^*$. We will consider the sum-product problem over finite fields. First, in $\mathbb{Z}$, if $A = \{1, \ldots, n\}$ is a geometric progression, we have $|A + A| = 2n - 1$. We also have $|A \cdot A| \gtrsim n^{2-\varepsilon}$. More specifically, $|A \cdot A| = \frac{n^2}{\log^\alpha n}$ for a constant α.

Theorem 11.1 (Gareu)**.**

$$|A + A|\,|A \cdot A| \geq c \cdot \min\left\{p\,|A|\,, \frac{|A|^4}{p}\right\}$$

In the range $\sqrt{p} << |A| << p$, this is a decent bound. We have a trivial bound

$$|A+A|\,|A\cdot A| \geq c\,|A|^2 \tag{11.1}$$

but the theorem in this range gives

$$|A+A|\,|A\cdot A| \geq |A|^{2+c} \tag{11.2}$$

and if $|A| \approx p^{2/3}$, then

$$|A+A|\,|A\cdot A| \geq c\,|A|^{5/2} \tag{11.3}$$

Now, back to the sum-product graph. Recall that we have $(p-1)p$ vertices (a,b) for $a \in \mathbb{F}_p^*$ and $b \in \mathbb{F}_p$. They are connected by an edge iff $ac = b+d$. Let M denote the adjacency matrix of G. One sad thing is that it has loops, a vertex can be connected to itself. We have $p(p-1)$ eigenvalues

$$\mu_0 \leq \mu_1 \leq ... \leq \mu_{p^2-p-1} \tag{11.4}$$

and let λ be the second largest eigenvalue. For any two vertices u, v, if $a \neq c$ and $b \neq d$, then they have exactly one common neighbour. To see this, let (x,y) be a common neighbour. We have $ax = b+y$ and $cx = d+y$. From this system of equations, we get $x = (b-d)(a-c)^{-1}$ and $y = x(a+c)-b-d$, which determines the point exactly. On the other hand, if $a = c$ or $b = d$, then there are no common neighbours. We have

$$M^2 = \mathcal{J} + (p-2)I - E \tag{11.5}$$

where E is an error matrix. The error term defines a graph G_E, with $(v_i, v_j) \in E(G_E)$ iff $a = c$ or $b = d$. The error graph is a regular graph, so all the eigenvalues are bounded in degree. G_{p-1} is a $(p-1)$-regular graph, so $p-1$ is the largest eigenvalue, and $\vec{1}$ is the eigenvector. There is an orthonormal basis consisting of the eigenvectors of M. Let θ denote the second largest eigenvalue of M. We know $|\theta| < p-1$. Multiply equation (10.5) by v_θ, and we get

$$(\theta^2 - p + 2)v_\theta = Ev_\theta \tag{11.6}$$

which means E has the same eigenvectors as M. So G_E is a $(2q-3)$-regular graph, so any eigenvalue of E is at most $2q-3$. We then obtain $\theta^2 - p + 2 \leq 2p$, so $|\theta| < \sqrt{3p}$. Now we use the spectral bound

$$\left| e(S,T) - \frac{|S|\,|T|}{p} \right| \leq \theta\sqrt{|S|\,|T|} \tag{11.7}$$

Defining the two sets, we choose $S = (A \cdot A) \times (-A)$, and $T = (A^{-1}) \times (A+A)$. There is an edge between $(a \cdot b, -c)$ and $(b^{-1}, a+c)$. The sizes satisfy $|S| = |A \cdot A|\,|A|$ and $|T| = |A| \cdot |A+A|$, and the number of edges in $e(S,T)$ is at least $|A|^3$. Thus,

$$|A|^3 \leq e(S,T) \leq \frac{|S|\,|T|}{p} + \sqrt{3p\,|S|\,|T|} = |A\cdot A|\,|A+A|\,|A|^2 + |A|\,\sqrt{3p\,|A\cdot A|\,|A+A|}$$

looking at each term in this sum individually, we get the result.

Lecture 12

Recall that G_n is dense if $e(G_n) \geq cn^2$ for $c > 0$ fixed.

Lemma 12.1 (Szemerédi's Regularity)**.** *If G_n is dense, then you can partition the vertices into subsets, where the edges between the 'most of' the subsets are 'randomlike' bipartite. In other words:*

For any $\varepsilon > 0$. there is an m such that any graph G_n for $n \geq n_0(\varepsilon)$ can have its vertex set partitioned into $\sim \frac{n}{m}$ sized partition classes so that for all but εm^2 pairs of the classes, the bipartite graphs between them are ε-regular.

Definition 12.2. *Let $G(A, B)$ be bipartite, $|A| = |B| = n$. Then $G(A, B)$ is ε-regular if for any $A' \subseteq A$ and $B' \subseteq B$, with $|A'| \geq \varepsilon |A|$, $|B'| \geq \varepsilon |B|$, we have*

$$\left| \frac{e(A', B')}{|A'|\,|B'|} - \frac{e(A, B)}{|A|\,|B|} \right| \leq \varepsilon$$

This quantity is the 'density'. We could write it as

$$|\delta(G(A', B')) - \delta(G(A, B))| \leq \varepsilon \tag{12.1}$$

Lemma 12.3. *If G_n is a graph where every edge is the edge of exactly one triangle, then G_n is sparse. However, you can get really close to n^2 edges. There are graphs with $\geq \frac{n^2}{e^{\sqrt{\log(n)}}}$*

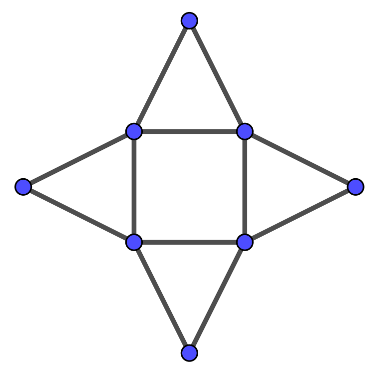

Now, a question. Consider $[n] \times [n]$, the integer grid. if $S \subset [n] \times [n]$, and $|S| \geq cn^2$, then S contains a "corner".

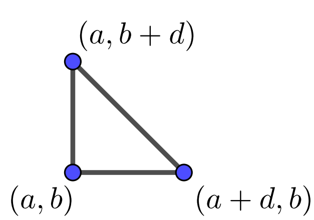

$d \neq 0$

This is a similar question to Roth's theorem, which is that if $S \subset [n]$ and $|S| \geq cn$, then S contains an arithmetic progression of length 3. However, the above question is stronger. You can see this with a projection trick.

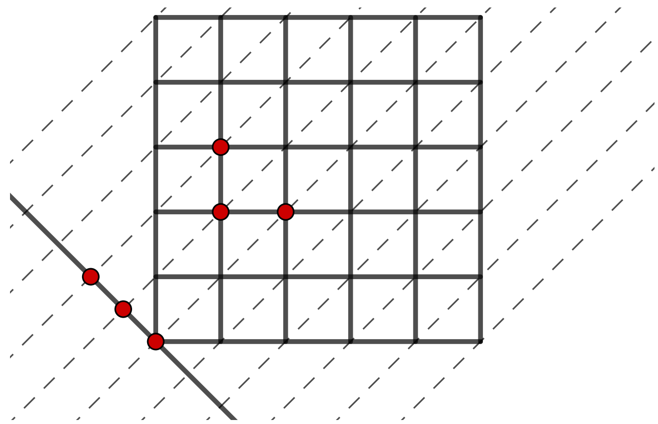

We can solve the question by creating three vertex sets. The vertices are the slope -1 lines in the grid, the horizontal lines, and the vertical lines. These sets have $2n$, n, and n vertices respectively. We join two vertices if their point of intersection is in S. Every point in S has exactly 3 lines going through it, so that is a triangle. If we assume this graph is dense, there is one edge sitting on two triangles. This gives us a corner. Otherwise it is sparse, but this contradicts density of S.

Lecture 13

Recall that a ε-regular pair is two sets, where the bipartite graph $G(V,W)$ satisfies the following condition. If $V' \subseteq V$, $W' \subseteq W$ and $|V'| \geq \varepsilon |V|$, $|W'| \geq \varepsilon |W|$, then

$$|\text{the density of } G(V',W') - \text{the density of } G(V,W)| \leq \varepsilon$$

Lemma 13.1 (Regularity, by Szemerédi in the 70's)**.** *For any $\varepsilon > 0$ there is an $M \in \mathbb{N}$ so that any graph G_n for n 'large enough' can be partitioned into M partition classes such that all but at most εM^2 partition class pairs are ε-regular. One can require that all classes have size $\in [\frac{n}{M} \pm 1]$.*

Given G_n and an ε-regular partition $V_1, \dots, V_M$, one can define the δ-reduced subgraph of G_n. This is called $G_n^* \subset G_n$, made by removing all edges between edges within one of the V_i, removing all edges between ε-irregular pairs, and removing all edges between pairs with edge density $< \delta$.

We had cn^2 edges originally, how many edges did we remove in this process? In the first step, at worst every partition graph contains a complete graph, so we could have lost up to

$$M \binom{\frac{n}{M}}{2} \approx M \left(\frac{n}{M}\right)^2 = \frac{n^2}{M}$$

In the second step, from the lemma, we could have removed ε complete bipartite graphs, and lost

$$\varepsilon M^2 \left(\frac{n}{M}\right)^2 = \varepsilon n^2$$

from the third step, we removed

$$\delta \left(\frac{n}{M}\right)^2 \binom{M}{2} \approx \delta n^2$$

therefore we have $e(G_n^*) \geq \left(c - \frac{1}{M} - \varepsilon - \delta\right) n^2$ edges in the new graph. So, if for example we chose our constants so that $\frac{1}{M} + \varepsilon + \delta < \frac{c}{2}$, then $e(G_n^*) \geq \frac{c}{2} n^2$ which is a significant number of edges. We can use this to prove the strange triangle lemma from last class.

Proof of 11.3. Another way to state the lemma is that for any $c > 0$, there is an $n_0(= n_0(c))$ so that if G_n is as in the lemma conditions, then $e(G_n) < cn^2$.

Suppose there is a $c > 0$ so that there are graphs G_n with arbitrary large n so that every edge is an edge in exactly one triangle and $e(G_n) \geq cn^2$. Now, use regularity lemma with ε small enough on G_n. Find the δ-reduced subgraph of G_n. Because every edge is an edge of exactly one triangle, if the number of edges we remove is less than the number of triangles, we will still have a triangle in the end. The number of triangles is $\frac{cn^2}{3}$. So if $\frac{1}{M} + \varepsilon + \delta < \frac{c}{3}$, then we will be left with a triangle.

This triangle connects three classes. We'll suppose they have size N vertices each. Then there are $3N$ vertices total, and δN^2 edges, and the three pairs are ε-regular and δ-dense. How many points are in V_2 with νN neighbours in V_1? If the number of these bad points is $> \varepsilon N$, we get a contradiction. That is, the density of $(BAD, V_1) > \delta - \varepsilon$ if $|BAD| \geq \varepsilon N$, so there not many vertices with small degree. This implies there are vertices in V_2 which are good for each of V_1 and V_3. Considering the neighbours of this point, the neighbouring sets have many edges between them. $\square$

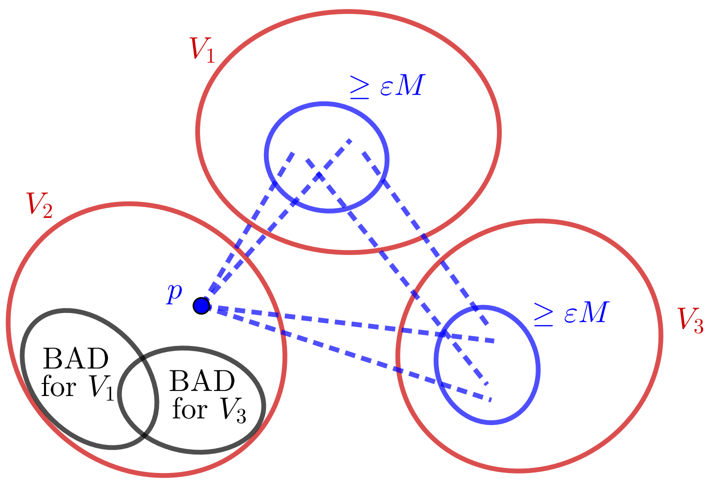

Lecture 14

H.W.4 is going to be a 1–2 page "essay"/"reflection" on the presentation, to be completed after the presentation is done. It should be about the main ideas of the proof, what you learned while reading the material and preparing the talk.

14.1 Regularity Lemma and Hypergraphs

An induced matching is a set of vertex-disjoint edges in G_n, where the vertices of the matching induced the edges of the matching only. So for a graph in general, G^* is an induced subgraph if there is a vertex set $V' \subset V(G_n)$ such that $V(G^*) = V'$ and $G^* \subseteq G_n$ and G^* contains all edges in G_n between vertices in V'.

Corollary 14.1. *If G_n is the union of n induced matchings then it is sparse.*

For simplicity, consider a bipartite graph G_n.

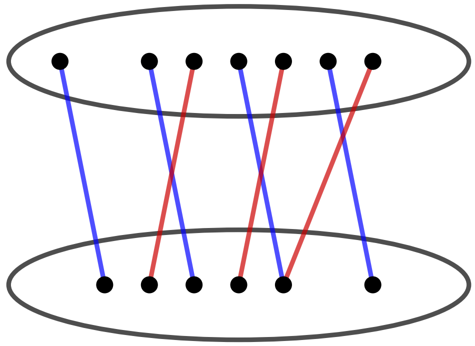

Proof sketch. Take M partition classes in an ε-regular partitioning. Let us consider the reduced (ε, δ) regular subgraph of G_n, G_n^*. Now, there is a matching M with $\geq c'n$ edges. If in a class V_i the number of vertices of $M < \varepsilon^* |V_i|$, then remove those edges from M. At the end, we removed at most $\varepsilon^* n$ edges from M.

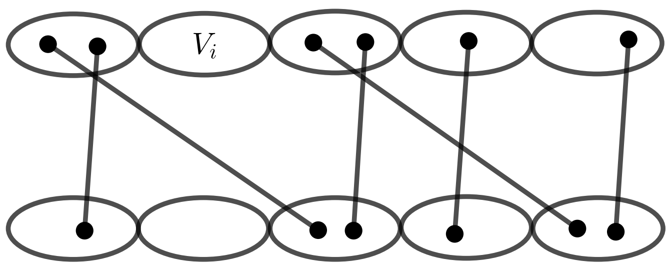

If $\varepsilon^* < c'$, there is some edge from M remaining. Say this edge joins classes V_i to W_j, so these classes have an edge which survived all the cleanings. This means there were $\geq \varepsilon^* |V_i|$ vertices from M in V_i and $\geq \varepsilon^* |W_j|$ vertices from M in W_j as well. Choose ε small enough so that $\varepsilon^* > \varepsilon$. Since the pair V_i and W_j are ε-regular and δ-dense, we can find a subgraph M_i of V_i and M_j of V_j, and the number of edges between them is $\geq \varepsilon^* |V_i| \varepsilon^* |W_j| \cdot \delta$. This means that there are many edges here, but since these edges came from n matching, so the max number of edges is $\min(|M_i|, |M_j|)$. □

Definition 14.2. *A 3-uniform hypergraph $\mathcal{H}_n^{(3)}$. is a graph where every (hyper) edge connects 3 vertices.*

Observe that the maximum number of edges in $\mathcal{H}_n^{(3)}$ is $\binom{n}{3} \approx n^3$. A complete hypergraph is usually called a clique. A modification of 11.3: if in $\mathcal{H}_n^{(3)}$ every edge is the edge of exactly one clique, then it is sparse. That is, for every $c > 0$, there is an $n_0 = n_0(c)$ so that if $\mathcal{H}_n^{(3)}$ has this property, and $n \geq n_0$, then the number of edges in this graph is $< cn^3$.

Lecture 15

Theorem 15.1. *If in $\mathcal{H}_n^{(k)}$, every edge is the edge of exactly one clique, then the graph is sparse (that is, $e(\mathcal{H}_n^{(k)}) = 0(n^k)$).*

This is a more general version of what was stated last time. If $k = 3$, then the above theorem implies: For each $c > 0$, there is an $n_0 = n_0(c)$ such that if $S \subset [n] \times [n]$ and $n \geq n_0$ and $|S| \geq cn^2$, then S contains a square.

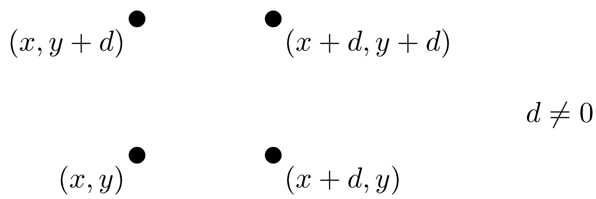

Like the previous theorem of this kind, we solve this with a projection trick, this time projecting a three dimensional set of points onto the two dimensional set. Let's assume S contains no square.

Then $\tau := S \times [n] \subset [n] \times [n] \times [n] \subset \mathbb{R}^3$ contains no configuration like (x, y, z), $(x+d, y, z)$, $(x+d, y+d, z)$, $(x, y+d, z+d)$. Points like this determine a tetrahedron. So we create planes slicing through the space.

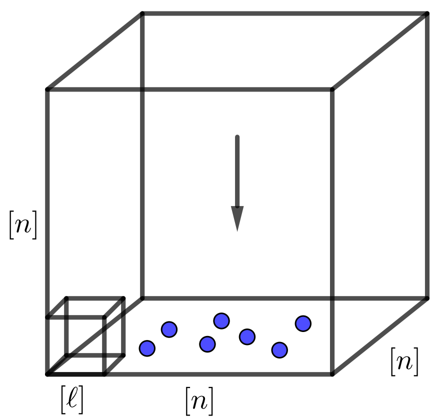

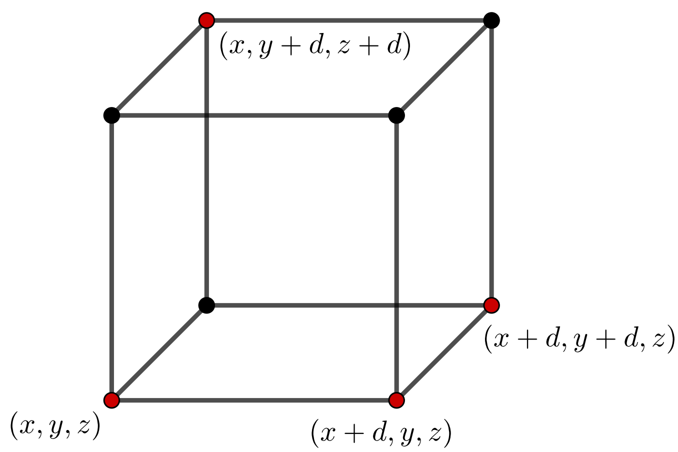

Choose a small number ℓ, and consider the smaller cube if this size within the real one. If three points p_1, p_2, $p_3 \in [\ell]^3$ determine a plane P, the number of planes parallel to P with non-empty intersection, with $[n]^3$ is $\leq c_\ell n$. This can be seen with some linear algebra. In fact, you can compute that $c_\ell = 3\ell$ works.

Define vertices of a graph to be one of the four plane edges that can make one of these tetrahedra. We connect three edges if the intersection of the planes is in τ. Each point of τ defines 4 edges, and every edge is an edge of a clique. If an edge is in at least 2 cliques, then one of the cliques defines a tetrahedron with sides parallel to the configuration we are looking for. This completes the proof.

It's possible to do a similar argument to show that there is always a full 3×3 grid in such a set. You project to the space $S \times [n]^7$, (in general, make the dimension one less than the number of points in your grid) and look at the hyperplanes.

Lecture 16 – Probabilistic Method

$G_{n,p}$ is a graph on n vertices with each edge having independent probability p of appearing. A famous example is $G_{n,1/2}$, every edge is taken independently with probability $1/2$. Question: what is the probability that $G_{n,1/2}$ is not connected? To start, create a partition into vertex sets m and $n-m$. There are $m(n-m)$ potential edges between them, so there is a $2^{-m(n-m)}$ chance these components are disconnected. Summing over all such partitions, we have a trivial upper bound on the probability that the graph will be disconnected.

$$\sum_{m=1}^{\lfloor \frac{n}{2} \rfloor} \binom{n}{m} 2^{-m(n-m)} \tag{16.1}$$

We can evaluate this with some cool splitting of the sum. Let $\varepsilon > 0$, and then

$$\sum_{m=1}^{\lfloor \frac{n}{2} \rfloor} \binom{n}{m} 2^{-m(n-m)} = \sum_{m=1}^{\lfloor \varepsilon n \rfloor} \binom{n}{m} 2^{-m(n-m)} + \sum_{m=\lfloor \varepsilon n \rfloor + 1}^{\lfloor \frac{n}{2} \rfloor} \binom{n}{m} 2^{-m(n-m)} \tag{16.2}$$

and estimate each term. The second one is easier, we have $m(n-m) \geq \varepsilon(1-\varepsilon)n^2$ so this sum is less than a sum of $\frac{1}{2^{cn^2-n}}$. It sounds like we should actually use $\varepsilon = \log n$?

Recall the Ramsay number $R(k,k)$, which is the least number such that any two colouring of the complete graph on $R(k,k)$ vertices results in a monochromatic complete subgraph of size at least k. We got the bound of 4^k for this number before. We will show that the lower bound is $\sqrt{2}^k$. In K_n, colour the edges independently with probability $1/2$. What is the probability that you have k vertices that span a monochromatic complete subgraph?

Given k vertices, there is a $\left(\frac{1}{2}\right)^{\binom{k}{2}-1}$ chance it will be monochromatic. There are $\binom{n}{k}$ k-tuples of vertices. If we sum over all these probabilities and the result is less than one, there might not be a monochromatic subgraph. It's pretty intuitive, but this uses a probability theorem called linearity of expectations.

Let $X = 1$ if monochromatic, 0 otherwise. Then consider the expected value of $X_{\binom{n}{k}}$. By linearity of expectations,

$$EXP\left(\sum_{S \subset V(G_n), |S|=k} X_S\right) = \sum_{S \subset V(G_n), |S|=k} EXP(X_S)$$

so if this is less than 1, there is a colouring without a monochromatic subgraph. Thus if we can show $\binom{n}{k} \cdot \left(\frac{1}{2}\right)^{\binom{k}{2}-1} < 1$ we are done. This inequality does hold if $n \sim \sqrt{2}^k$.

Lecture 17 – Extremal Set Theory

We focus on finite sets. S_n is an n element set, and $\mathcal{F} \subset 2^{S_n}$ is a family of subsets of S_n. A typical question in Extremal Set Theory is 'What is the maximum $|\mathcal{F}|$ under some given conditions?' You could require the subsets to be 3 elements each for example. The extremal set is then a 3-regular hypergraph.

Here's a complicated example. Let $\mathcal{F}$ have the property that if A, $B \in \mathcal{F}$, then for all $A \neq B$, $A \cap B = k$ for some fixed integer $k \in \mathbb{N}$. What is $\max |\mathcal{F}|$ under this condition?

We can assume that for all $A \in \mathcal{F}$, $|A| > k$ otherwise if there was a set with $|A| = k$ every pair of sets B, $C \in \mathcal{F}$ not equal to A have to have $B \cap C = A$. We then have $B - A$ and $C - A$ are disjoint, and every $B - A$ set has at least one point outside A. This gives us at most $n - k + 1$ sets in this case.

Now, every $A \in \mathcal{F}$ can be represented by its n-dimensional indicator vector v_A. We can then consider the Gramm matrix of the indicator vectors i.e. the matrix where entry a_{ij} is the dot product of vector i with vector j. The dot product of two indicator vectors is the size of the intersection. This matrix will thus be all k's, except along the diagonal where there will be the cardinalities of the sets. Let $m = |\mathcal{F}|$.

$$M_{\mathcal{F}} = \begin{pmatrix} |A_i| & k & ... & k \\ k & |A_2| & ... & k \\ ... & ... & ... & ... \\ k & k & ... & |A_m| \end{pmatrix} \tag{17.1}$$

We want to bound the rank of this matrix in terms of m, k, and n. We can use the bound $Rank(M_1 + M_2) \leq Rank(M_1) + Rank(M_2)$. It follows that $Rank(M_{\mathcal{F}}) \geq m - 1$, which you can see by how $M_{\mathcal{F}} - k\mathcal{J}$ is diagonal with nonzero entries in the diagonal, and

$$m = Rank(M_{\mathcal{F}} - k\mathcal{J}) \leq Rank(M_{\mathcal{F}}) + Rank(k\mathcal{J}) \tag{17.2}$$

Now, the definition of this matrix is via $\langle v_A, v_B \rangle = x_1y_1 + ... + x_ny_n$, so we can split our matrix up into $M_{\mathcal{F},1} + M_{\mathcal{F},2} + ... + M_{\mathcal{F},n}$ where $M_{\mathcal{F},j}$ is the j^{th} coordinate of the dot product, x_jy_j. Every column is a constant multiplier of the first column, so every $M_{\mathcal{F},j}$ has rank 1. Thus

$$Rank(M_{\mathcal{F}}) \leq \sum_{j=1}^{n} Rank(M_{\mathcal{F},j}) = n \tag{17.3}$$

We thus have $m - 1 \leq Rank(M) \leq n$ so a bound is that $m \leq n + 1$.

Theorem 17.1. *Consider $\mathcal{F}$ where for all $A, B \in \mathcal{F}$, $A \not\subset B$. Then* $\max |\mathcal{F}| \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}$.

Proof. Let us suppose that elements of S_n are labelled and let's consider permutations of the n elements. A permutation $\pi(n)$ is "good" if there is a slice given by $A \subset \mathcal{F}$; that is, the first $|A|$ elements of $\pi(n)$ are exactly the elements of A.

If $\pi(n)$ is good for both A and B, then $A = B$ due to the restriction on $\mathcal{F}$. There are exactly $|A|!\,(n - |A|)!$ permutations which are good for a particular $|A|$ because the order of the first $|A|$ elements and the last $|n - |A||$ elements of $\pi(n)$ don't matter. Now,

$$\begin{aligned} &\sum_{A \in \mathcal{F}} |A|!\,(n - |A|)! \leq n! \\ &\implies \sum_{A \in \mathcal{F}} \binom{n}{|A|}^{-1} \leq 1 \end{aligned} \tag{17.4}$$

this last equation is called the LYM inequality. Every entry in this sum is $\geq \binom{n}{\lfloor \frac{n}{2} \rfloor}^{-1}$, so

$$|\mathcal{F}| \binom{n}{\lfloor \frac{n}{2} \rfloor}^{-1} \leq \sum_{A \in \mathcal{F}} \binom{n}{|A|}^{-1} \leq 1 \tag{17.5}$$

□

Lecture 18 – Ramsey Theory

We studied the Ramsey number $\sqrt{2}^k \leq R(k,k) \leq 4^k$ in this class. There is an analogue of this for Hypergraphs as well. A large part of this field is *Arithmetic Ramsey*, which includes Shur's Theorem, and results like the density of a subset of $\mathbb{N}$ being large means there are arithmetic progressions of any size.

18.1 Geometric Ramsey

Theorem 18.1 (Erdős-Szekeres)**.** *Any sequence of real numbers has a 'long' increasing subsequence or a 'long' decreasing subsequence (not necessarily* strictly *increasing or decreasing).*

For example, the sequence $1, 2, 1, 1, 3, 4, 5, 3$ has the increasing subsequence $1, 1, 1, 3, 4, 5$ and decreasing $2, 1, 1$. Now, what we mean by long is 'size at least $\sqrt{n}$'. The exact statement in the theorem is that $s \cdot r + 1 \geq n$ where s is the longest increasing, r is the longest decreasing. You can see that there is a construction with both sequence lengths at most $\sqrt{n}$, just take $1, 2, ... \sqrt{n}, 1, ...$

Proof. Let $a_i \in \mathbb{R}$ for $1 \leq i \leq n$. label each element in the sequence. The labelling of element a_i is (b_i, c_i), where b_i is the size of the largest increasing sequence in the first i elements ending in a_i, and c_i is the size of the largest decreasing sequence in the first i elements ending in a_i. Observe that every label is distinct. To see this, note that we have two cases for elements a_i, a_j. If $a_i \geq a_j$, then $c_i < c_j$. If $a_i < a_j$, then $b_i < b_j$.

Now, if s is the size of the longest increasing sequence, r is the longest decreasing sequence, then all of our b_i are in $[s]$ and all of our $c_i \in [r]$. There are at most sr pairs (b_i, c_i), so $sr \geq n$. □

Note that unlike the infinite arithmetic progression problem, you can always find an infinite increasing or decreasing sequence, not just one of arbitrarily large size. You just let S be the elements a_k in our set for which every element after a_k is smaller. If S is infinite, it is an infinite decreasing subsequence. If its finite, then starting at the element after the final element of S, you can start find an infinite increasing sequence.

Theorem 18.2 (Kőnig)**.** *An infinite rooted tree with finite degree vertices has an infinite path.*

Another famous theorem in this area is the following.

Theorem 18.3 (Erdős-Szekeres Happy Ending)**.** *Given n points in the plane, no three in one line, then there are $f(n)$ points in convex position. Let $F(n)$ = least number so that among $F(n)$ points, there is at least n in convex position. The theorem states $2^n \leq F(n) \leq 4^n$.*

It's not hard to see that $F(4) = 5$, that is, for any 5 there are 4 in convex position. To see this, assume that the convex hull was not already a quadrilateral, so two points are inside the hull. Connect the two with a line. This line hits two sides of the triangle, and the 2 points determining the third side are in convex position with the 2 interior points.

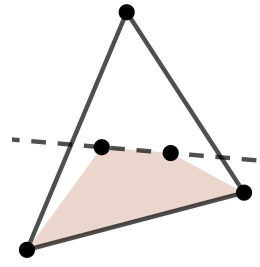

One more application; let $B(k,\ell)$ be the least number so that any point set of size $\geq B(km,\ell)$ contains a k-cAp or an ℓ-cUp. These things are points arranged in some kind of parabolic position. We can prove by induction that

$$B(k,\ell) \leq \binom{k+\ell+4}{\ell-2} + 1 \tag{18.1}$$

Lecture 19

Book recommendation: Lovász, Problems and Exercises in Combinatorics is a book you *must have* if you're into combinatorics.

Let C_t^d be a 'combinatorial space' of dimension d over t characters. Words of length d composed using t characters. The 'characters' could be anything, but you can take for example $t = \{0,1,2,\ldots,t-1\}$. A word could be $[0,5,2,1,0,0]$ which is an element of C_6^6.

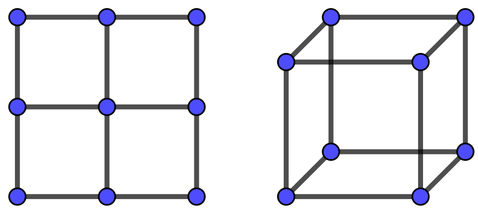

C_3^2 C_2^3

A combinatorial line in C_t^d consists of t words (points) such that in position i for $1 \leq i \leq d$, one of two cases holds. Either all entries are the same, or all entries are increasing from 0 to $t-1$. What this might look like is, with the first coordinate stationary and the second increasing,

$$\begin{matrix} (2, & 0, & \ldots &) \\ (2, & 1, & \ldots &) \\ (.. & .. & \ldots &) \\ (2, & t-1 & \ldots &) \end{matrix}$$

Consider the diagram below. In this space, you can count there are 7 combinatorial lines. You can see that the red line here is not a combinatorial line as it is decreasing in one coordinate, but the green one is.

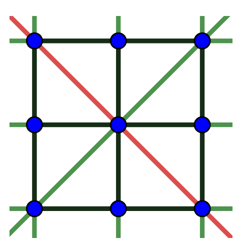

In more generality, note that $\left|C_t^d\right| = t^d$. We can see there are $(t+1)^d - t^d$ combinatorial lines. To see this, you can choose k coordinates to be constant, and the other $d-k$ will be increasing. This can be done in $\binom{d}{k}$ ways. Then there are k^t ways to choose which values are fixed at, so we have $\sum_{k=0}^{d-1} \binom{d}{k} k^t$ ways to do this.

Another argument is to extend the combinatorial space one character longer. Each combinatorial line will arrive at a unique point on the boundary of C^d_{t+1}, as everywhere you have $t+1$ must have been increasing, as everywhere else you have the same constant. That is, $C^d_{t+1} - C^d_t$ is in bijective correspondence with the lines in C^d_t.

Consider $C^d_2 =: Q^d$, the d-dimensional cube. A line in C^d_2 might look like

$$\begin{matrix} (0, & 1, & 1, & 0, & 0, & 1, & 0, & 0, & 0) \\ (1, & 1, & 1, & 0, & 1, & 1, & 0, & 1, & 0) \end{matrix}$$

Theorem 19.1 (Hales-Sewett)**.** *For any t and k there is a number $HJ(t,k)$ such that for any k-colouring of C^d_t there exists a monochromatic combinatorial line provided $d \geq HJ(t,k)$.*

We can prove this for C^d_2. Note that the points of C^d_2 represent the subsets of a d element set. We know that the maximum cardinality subset of subsets without one containing another is $\leq \binom{d}{\lfloor d/2 \rfloor} \sim \frac{2^d}{\sqrt{d}}$. This means that for and $c > 0$ there is a d_0 depending on c only so that if $d \geq d_0$, and we have a set $S \subset C^d_2$ with $|S| \geq c2^d$, then S contains a combinatorial line (which, again, is a set containing another). This proves it, as your colouring classes are eventually large enough so that one of them must contain a line.

We can also show this for $HJ(2,k)$. Well, we have chains of $d+1$ sets, and if any two have the same color, we have a line, so $HJ(2,k) = k$. In general, there is a way to use these as a base case for induction.

There is a harder 'density' version of Theorem 17.1, called the DHJ Theorem. The first proof was by Furstenberg and Katznelson, but there is a Polymath project proof of that as well. We will now state a couple applications of Theorem 17.1.

Theorem 19.2 (Vanderwaerden)**.** *For any k and r, there is a bound $W(k,r)$ such that any k colouring of the first n integers contains a monochromatic arithmetic progression of length r, provided $n \geq W(k,r)$.*

Theorem 19.3 (Euclidean Ramsey)**.** *For any point set $\{p_1, \dots p_r\} = P$, and any k-colouring of $\mathbb{R}^2$, there is a monochromatic scaled & translated copy of P.*

To begin to prove this, lets say our alphabet is $\{v_1 v_2, \dots, v_r\}$, and here $r = t$. If we have combinatorial line we can sum up the rows like so,

$$\begin{matrix} (v_i & \dots & v_1 & \dots & v_1 & \dots) \to V + mv_1 \\ (v_i & \dots & v_2 & \dots & v_2 & \dots) \to V + mv_2 \\ & & & \dots & & \\ (v_i & \dots & v_r & \dots & v_r & \dots) \to V + mv_r \end{matrix} \tag{19.1}$$

where m is the number of running coordinates, and V is the fixed sum of the fixed entries. Let the color of a word in the combinatorial space be the color at the spot on the plane that you land when you sum up the vectors in each line. A monochromatic combinatorial line will then give you a way to transform your set (based on V and m) to land on a monochromatic set of points.

Start with an $n^p \times n^p$ grid. This gives n^{2p} directions. Populate an $n^{\frac{1}{2}+c} \times n^{\frac{1}{2}+c}$ grid with points along those direction lines. There are at least $n^{\frac{1}{2}+c}/n^p = n^{1/2+c-p}$ points along each line, so there are $n^{\frac{1}{2}+c+p}$ points total. This means $\frac{1}{2} = c + p$. The number of projections onto each of those lines is n

A line contains $(a+a, b \cdot b)$ and $(c+c, d \cdot d)$. Direction vector is $(a+a-c-c, b \cdot b - d \cdot d)$. Label the n vertices with a and b, and suppose the average cell is about 1. We need that $a_j b_{n-j} \approx j(n-j)$.