

SUMMER REU 2009
APPRENTICE COURSE
PART I

LECTURES BY MIKLÓS ABÉRT

NOTES BY DAVID CHUDZICKI, AARON MARCUS, MIKE MILLER, JONATHAN STEVENSON
EDITED BY AARON MARCUS

CONTENTS

1. Class 1	1
2. Class 2	4
3. Class 3	8
4. Class 4	10
5. Class 5	15
6. Class 6	18
7. Class 7	22
8. Class 8	24
9. Class 9	28

1. CLASS 1

What is a Graph?

Definition 1.1. A *graph* is a set of vertices V and a set of edges E , where each edge is a two element subset of the set of vertices. A *directed graph* is a set of vertices V , and a set of edges E , where $E \subset V \times V$. (In both, loops and multiple edges are allowed.)

A graph is called *simple* if it is undirected, has no loops, and has no multiple edges. (For the first part of the course, we'll just worry about simple graphs.)

The *degree* of a vertex is the number of edges containing the vertex.

Example 1. Let V be the set of people in the classroom, and there is an edge between two people if the two people are friends. Is this graph directed? Alternatively, draw an edge between person 'A' and person 'B' if 'A' is attracted to 'B'. Is this directed?

Exercise 1. Is it possible that everyone in the class knows a different number of people? Here, we're assuming that the relation of "knowing each other" is symmetric, i.e, x knows y if and only if y knows x .

Solution 1. It is impossible: say there are n people in the class. Then, (since one can't know oneself) the number of people that each person could know is one of $0, 1, \dots, n - 1$.

Date: July 23, 2009.

So, for each person to know a different number of people, there must be some person that knows everyone, and there must be some person that knows no one. But this is impossible: the relation of knowing each other is symmetric, so the person who knows no one should know the person who knows everyone!

So, we just proved that there is no finite (simple) graph such that every vertex has a different degree.

Problem 1. Is there an infinite simple graph such that all degrees are different?

Guess 1. Of the people who voted, “it exists” was the consensus, but there wasn’t a very high turnout!

Exercise 2. There are six hippos sitting in a sauna. Show that either

- (1) there are three of them that all know each other, or
- (2) there are three of them that all do not know each other.

Also, show that this fails if there are only five hippos in the sauna.

Solution 2. We can assume that there’s a vertex v with degree at least three, by considering the complement if necessary. Now, look at least three of the vertices that are connected to v , say v_1, v_2 , and v_3 . If none of v_1, v_2 , and v_3 know each other, then they form an empty triangle, and if two of them know each other, say v_1 and v_2 , then v, v_1 , and v_2 form a full triangle.

What does this last exercise mean in terms of graphs? It means that in any (undirected) graph with six points, you can either find a (full) triangle or an ‘empty’ triangle.

What do we mean when we say two graphs are the same?

Definition 1.2. We say that two graphs are *isomorphic* if there is a bijection between the vertex sets that preserve the edges.

Exercise 3. Prove that all counterexamples to the “hippo theorem” with five points are isomorphic to the pentagon.

Proof. The degree of every vertex must be two! □

Exercise 4. Suppose that G is a simple graph on n points, where the degree of every vertex is 2. Show that G is a disjoint union of cycles.

Proof. Start at a vertex, go to a neighbor, go to the next (unused neighbor), keep going until \dots , we got back where we started! If you used all the points, you’re done, if not, keep going! □

Now, let’s talk about prom. We’ll get to the first real theorem in graph theory. We need to get all the boys a date to prom. Let B be the set of boys and G the set of girls. We’ll assume that there are more girls than boys. We’ll draw a graph where the set of vertices is the set of boys and girls, and we draw an edge between a boy and a girl if the boy is willing to dance with the girl. (This is an example of a *bipartite* graph, because boys don’t want to dance with boys and girls don’t want to dance with girls. In other words, there are two big potatoes, and the edges only go between the potatoes; there are no edges internal to each potato.)

Now, the question is, “Can we arrange things so that every boy is dancing with a girl with whom he wants to dance?” The following condition is necessary: we *cannot* find a subset of k boys that will only dance with less than k girls. In fact, this condition is sufficient.

Theorem 1.1. In the above situation, if for all $X \subset B$, X has at least $|X|$ neighbors in G , then there is a matching from B to G . (In other words, for every boy, we can find a unique neighbor.)

Proof. Start matching people in a “really stupid way”. If we can do it forever great! If not, we’ll get to a point where we have a bijection between a set of boys B_0 and a set of girls G_0 , and we’re trying to match a boy $b \notin B_0$, and we can’t, because every girl b wants to dance with is in G_0 . Then, consider the set of girls G' that consists of the girls that can be reached from b by an “alternating path”. If G' is a subset of G_0 , let B' be the set of boys in B_0 matched to girls in G' . Then, the set of neighbors of $B' \cup \{b\}$ is G' , which contradicts the hypothesis of the theorem. So, G' is not a subset of G_0 , and we can re-arrange the partnerships along an alternating path that ends outside of G_0 . This allows us to continue our matching. \square

Definition 1.3. We say that a graph G is k -regular if the degree of every vertex is k .

Corollary 1.1.1. Every k -regular ($k \geq 1$) bipartite graph has a perfect matching. (So in the language above, every boy will dance with exactly k girls and for every girl, there are exactly k boys who want to dance with her.) Why does the condition imply that there are the same number of boys and girls?

Proof. Apply the marriage lemma: this gives a matching; it must be perfect because there are the same number of boys and girls!

(To see that there are the same number of boys and girls, note that if there are m boys and n girls, then there must be $k \cdot m$ edges in the graph, since each boy will dance with k girls. However, there are also $k \cdot n$ edges in the graph, since for every girl, there are k boys that will dance with her. Thus, $k \cdot n = k \cdot m$, i.e, $m = n$.)

(To see that the lemma applies: let X be a subset of the boys. Let X be any subset of the boys, say it has n elements. Then, call Y the set of neighbors of X . Then, there are $n \cdot k$ edges going from X to Y . If there are less than n elements in Y , then at least one of the points of Y has degree greater than k , which is impossible.) \square

Exercise 5. Shuffle a deck of (fifty-two) cards, and deal thirteen hands of four cards each. Show that you can choose one card from each hand so that you obtain exactly one card of each rank (i.e., one ace, one two, one three, etc.).

Proof. Use the above corollary: the group of boys is each of the hands, and the girls are the ranks of the cards. The cards are the edges. This is a 4-regular graph, so we’re done.

There is a slight problem: there could be multiple edges, but everything (the marriage lemma and the corollary) works with multiple edges. \square

Problem 2. Suppose we have an $a \times b$ “table of cards”. A *horizontal shuffle* is when we permute the elements of the first row, then permute the elements of the second row, etc. A *vertical shuffle* is defined with columns in place of rows. Show that one can obtain any permutation of the table by performing a vertical shuffle, then a horizontal shuffle, and then a vertical shuffle.

Hint: first solve it for a 2×2 grid.

Question from the class: is it true in three dimensions? Answer: Probably, if you do the first dimension, then second, then third, then second then first.

Definition 1.4. Let G be a graph. A *path* in G is a sequence of vertices x_0, x_1, \dots, x_d such that x_j is a neighbor of x_{j+1} for all $0 \leq j < d$.

We say that $x \sim y$ if there is a path from x to y . Is \sim an equivalence relation? Yes, if G is undirected, but if G is directed, then maybe not. Equivalence classes are called the *connected components* of G . In particular, we say that G is *connected* if it has one component. In other words, G is connected if between any two points there is a path.

How many edges must a connected graph have?

Lemma 1.1. If G is a connected graph of n points, then it must have at least $n - 1$ vertices.

Proof. To prove this seemingly innocuous fact, we need some preliminaries.

Definition 1.5. A *tree* is a connected simple graph with no cycles.

Claim 1. A tree with n vertices has exactly $n - 1$ vertices.

Proof. We need to find a leaf, since then we can delete that leaf and use induction (the tree is still connected, and the number of vertices and edges each drop by one). Why does every finite tree have a leaf? If there were no leaf, we'd have a cycle: start at any vertex, and then follow a path, never choosing the same vertex as before. We have to get back to some point we already used! \square

Now, we can prove the lemma. We should prove that every connected graph contains a tree as a subgraph, since then the claim finishes our lemma. For each cycle, remove an edge, and after a while, there will be no edges left. \square

A related question: does every connected graph (not necessarily finite) contain a tree?

A question for next time: how many trees are there on n vertices? We'll do this by first solving it for $n = 1$, then for $n = 2$, then for $n = 3$, etc. There is 1 tree for $n = 2$, and there are 3 trees for $n = 3$, 16 trees for $n = 4$, and 125 trees for $n = 5$. We conjecture that there are exactly n^{n-2} trees on n points.

2. CLASS 2

Corollary 2.0.2. Every k -regular bipartite graph is the disjoint union of k perfect matchings.

Proof. It has a perfect matching, so remove it and induct on k . \square

Definition 2.1. A *weighted graph* is a graph equipped with a positive weight for every edge. A weighted graph is *regular*, if for every vertex, the sum of weights of edges leaving the vertex is a constant.

Lemma 2.1. Is the following conjecture true or not true? Every regular weighted bipartite graph has a perfect matching.

Proof. We can certainly do it if all weights are one, since this is the original marriage lemma. If every weight is an integer, it's easy, since we can replace a single edge with a weight of n with n edges all with weight one. Then, we're in the situation of every edge having weight one. Now, we can do all rational weights, since we can clear denominators and return to the situation of integers. We'd like to finish by using the density of the rationals in the reals (and approximating our graph with real weights by graphs with rational weights), but problems occur because we cannot ensure that the "approximating graphs" are regular. A proof exists (using these methods), but we'll need to use linear algebra.

For another approach, we'll think about what the marriage lemma actually says for regular bipartite graphs. It says we have a perfect matching if and only if for every subset of size n that subset has at least n neighbors? (How do we know the matching is perfect? In other words, how do we know that each side has the same number of points? Use the same proof as for k -regular graphs: add up all the weights on the left side, you get $n \cdot c$, where n is the number of points and c is the constant value of the weights at each vertex. On the right, we have $m \cdot c$. But these must be equal, so we have $n \cdot c = m \cdot c$, i.e., $n = m$.)

To prove the marriage lemma for regular graphs, we just need to check the condition from the marriage lemma. But, the condition from the lemma follows immediately from the regularity condition: if X is a subset of the left hand side and Y is its set of neighbors, we can conclude that $c|X| \leq c|Y|$, i.e., $|X| \leq |Y|$. \square

(From the marriage lemma for regular graphs, we get a similar corollary for regular graphs as the first corollary from today.)

Exercise 6. Euler's Königsberg's Bridges Problem. One can't take a path that uses every bridge exactly once (see figure 1). It's really a graph theory question: add a vertex for each piece of land and an edge for each bridge, as in figure 2. We're asking if we can find a path that uses every edge exactly once.

How do we prove that no such path exists? Suppose you had a path: it starts somewhere and it ends somewhere. Now, for every point that is not a starting point or an ending point, one must visit that point an even number of times (since every time one visits it, one also leaves it). Moreover, one must leave on a different edge than the edge that one entered on. But this is impossible: every vertex has odd degree!

Definition 2.2. A graph is *Eulerian* if there is admits an *Eulerian circuit*, i.e., a path that uses every edge exactly once and ends where it starts.

For example, a cycle is Eulerian.

So, thinking about the Königsberg's Bridges Problem, we see that if a graph is Eulerian, then every vertex has even degree. In fact, the converse is true: if every vertex has even degree, then the graph is Eulerian. The full and precise statement is

Theorem 2.1. G is an Eulerian graph if and only if G is connected and every degree is even. (G must be finite, with no isolated points.)

Proof. We already saw the forward implication. For the reverse, we'll induct on the number of edges. Just start walking and eventually you'll get back to where you started. (You can always continue, because each vertex has degree 2.) If we created an Eulerian circuit, great! If not, remove that circuit. Then, the resulting graph is a disjoint union of connected graphs with all degrees even. By induction, each of these has an Eulerian circuit. We can combine

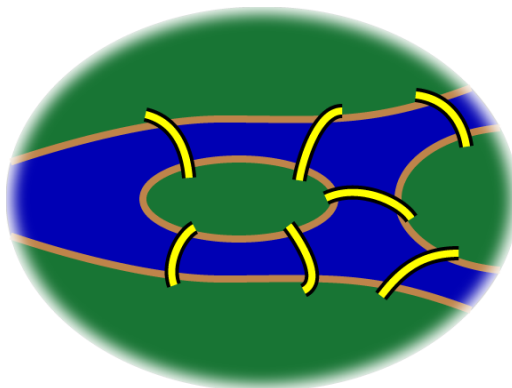


FIGURE 1. The Seven Bridges of Königsberg

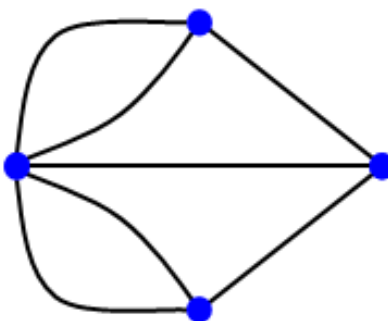


FIGURE 2. The same picture, translated into a graph

these Eulerian circuits together in the following manner. Start by taking the original path, but when you get to a vertex that intersects one of the remaining components, follow the Eulerian circuit on the component before resuming the path. In this way, we obtain a circuit for the entire graph. \square

Theorem 2.2. G is bipartite if and only if G has no cycle of odd length. (Recall that bipartite means that one can color the vertices with two colors such that no edge has two endpoints of the same color.)

Proof. The forward direction is immediate, since any graph with a cycle of odd length cannot be bipartite. For the reverse, just start somewhere and color it white; then color all its neighbors green (for Martian!); then color all the neighbors of all the green points white; keep going! There can't be any problems, because there are no odd cycles. \square

Now, we'll return to connectedness for directed graphs.

Definition 2.3. A directed graph is *weakly connected* if for all vertices x and y there is an undirected path between x and y . (In other words, the graph is connected when thought of as an undirected graph.)

A directed graph is *strongly connected* if for all vertices x and y there is a directed path from x to y . Note that strongly connected implies weakly connected but not conversely.

We say that a directed graph is *regular* if the “in degree” is equal to the “out degree”.

Exercise 7. Can you find a finite, regular, weakly connected directed graph that is not strongly connected?

NO! By mimicking the proof for Eulerian circuits of undirected graphs, but using regularity instead of evenness of degree, we see that the graph must have an Eulerian circuit, and is therefore strongly connected! So, we’ve proven the following fact:

Theorem 2.3. A finite, regular, weakly connected graph strongly connected.

Definition 2.4. A *permutation* on X is a

- bijection $X \rightarrow X$
- 1-regular directed graph on X
- rook configuration on X
- perfect matching between X and X
- cycle form

A *rook configuration* is an X by X matrix where we place rooks so that no two rooks are attacking each other. So, there is exactly one rook in each row and exactly one rook in each column. A rook configuration is also called a *permutation matrix*.

Now, we’ll give an example of cycle form. The cycle form of the permutation

$$1 \mapsto 3 \qquad 2 \mapsto 1 \qquad 3 \mapsto 2 \qquad 4 \mapsto 5 \qquad 5 \mapsto 4$$

is

$$(1, 3, 2)(4, 5)$$

How do we take the product of two cycles? What’s the product of $(1, 3, 2)$ and $(2, 1, 3)$? Uh Oh! In which order do we do the multiplication? We view the permutations as functions, so we do the one on the right first, then the one on the left as below.

$$(1, 3, 2)(2, 1, 3) = (1, 2, 3)$$

How do we take the product of two 1-regular directed graphs? We just put them next to each other.

How do we take the product of two rook configurations? We use the matrix product!

Definition 2.5. The *adjacency matrix* of a graph is a matrix where the rows and columns are labeled by the vertices of the graph, and the $v_i v_j$ entry is the number of edges going from v_i to v_j .

So, if a graph is undirected its adjacency matrix is symmetric. Loops are the entries on the main diagonal. The out degree of a vertex is the sum of the row for the vertex, the in degree of a vertex is the sum of a column for the vertex.

What does the adjacency matrix look like for a bipartite graph? It’s a block matrix:

$$\begin{pmatrix} 0 & A^t \\ A & 0 \end{pmatrix}.$$

So, for bipartite graphs, we form the *bipartite adjacency matrix*, which is just the matrix A , since this contains all the information as the “full adjacency matrix”. (More explicitly, the bipartite adjacency matrix is where we index the rows by one of the ‘potatoes’ and index the columns by the other ‘potato’.)

Problem 3. Every $2k$ -regular (undirected) graph contains a 2-regular graphs.

Problem 4. If M is an $n \times n$ matrix with non-negative entries and every row sum and every column sum is 1, then M is a convex combination of permutation matrices.

A convex combination of v_1, \dots, v_n is a sum

$$\lambda_1 v_1 + \dots + \lambda_n v_n,$$

where $0 \leq \lambda_i$ for all i , and $\lambda_1 + \dots + \lambda_n = 1$.

We haven't proved the first problem, but the second one we already proved today!

3. CLASS 3

Problem 5. (From last time.) Every $2k$ -regular (undirected) graph contains a 2-regular graphs.

Solution 3. Assume the graph G is connected (if not, we can do the following for each component). Since all vertices have even degree, there is an Eulerian circuit, which determines an orientation for each edge. Then there are k edges leaving/entering each vertex. We can double the graph, letting G_1 and G_2 be two copies, and let a new (bipartite) graph have one edge for each edge in the original graph, leaving the corresponding source vertex in G_1 and entering the corresponding target vertex in G_2 . This bipartite graph is k -regular, so has a perfect matching. This matching corresponds to a 2-regular subgraph in the original graph.

We can phrase the same argument using matrices: Orient the edges as above. Then we have the adjacency matrix of this directed graph, which has k 1's in each column (and is not symmetric). So this is also the adjacency matrix of a k -regular bipartite graph, and a perfect matching corresponds to a permutation submatrix, which for the original graph is a 2-regular subgraph.

Vector Spaces. Review of definitions: 'Vector space,' 'linear combination' (*finite* sums of $\sum \alpha_i v_i$), 'linearly independent', 'parallel' (one being a non-zero scalar multiple of the other).

Example 2. An infinite linearly independent set: $\{x^i : i \in \mathbb{N}\}$, in the vector space of polynomials (over some field).

Question: If a polynomial is zero as a function, does it have to be the zero polynomial? No. This never holds over finite fields since there are infinitely many polynomials but only finitely many functions. Subtract two different polynomials that give the same function to obtain a nonzero polynomial representing the zero function.

Degree: $\deg(fg) = \deg(f) + \deg(g)$. So we say $\deg(0) = -\infty$.

Example 3. In the space of functions from \mathbb{R} to \mathbb{R} ,

$$\{1, \sin(x), \sin(2x), \sin(3x), \dots, \cos(x), \cos(2x), \cos(3x), \dots\}$$

is a linearly independent set. Try to prove this, as it is an important example.

Definition 3.1. S is a maximal linearly independent set if S is linearly independent and for any T with $S \subsetneq T$, T is linearly dependent.

Lemma 3.1. If S is a maximal linearly independent set in V , then S spans V .

Proof. Let $v \in V$. Then $S \cup \{v\}$ is linearly dependent, so $\alpha v + \sum \alpha_i v_i = 0$ for some choice of $\alpha, \alpha_i \in F$, $v_i \in V$. Then $\alpha \neq 0$, or else we would have a linear dependence in S , so we can solve for v in terms of the v_i , so $v \in \text{span}(S)$. \square

Lemma 3.2. If $\{v_1, \dots, v_n\}$ is linearly independent, and $\{w_1, \dots, w_m\}$ spans V , then $n \leq m$.

Proof. Note: If S_1 spans S_2 , and S_2 spans S_3 , then S_1 spans S_3 . We will show that there is a w_i so $\{w_i, v_2, \dots, v_n\}$ is linearly independent, and a w_j so $\{w_i, w_j, v_3, \dots, v_n\}$ is linearly independent, and so on, until we have a list of n linearly independent w 's. Then if $n > m$, there is a repeated s , contradicting linear independence. For ease of notation, we only show the first step in this substitution argument. (The rest proceed by the same argument.) To replace v_1 by a w_i and keep linear independence, we need $w_i \notin \text{span}(v_2, \dots, v_n)$. There must be some w_i satisfying this, since if all w_i were in the span of $\{v_2, \dots, v_n\}$, then $\{v_2, \dots, v_n\}$ would span V , so we would have $v_1 \in \text{span}(\{v_2, \dots, v_n\})$, contradicting linear independence. \square

Lemma 3.3. If S_1 and S_2 are both linearly independent and spanning (we call such sets *bases*), then $|S_1| = |S_2|$.

Proof. We apply the previous lemma with S_1 linearly independent, and S_2 spanning, then again reversing the roles. \square

Definition 3.2. If S is basis for V , then $\dim(V) = |S|$, the dimension of V

Definition 3.3. The rank of S is the size of a maximal independent subset, or equivalently, the dimension of the span of S .

Definition 3.4. A subspace of a vector space is a nonempty subset closed under addition and scalar multiplication. Notation: $U \leq V$ means U is a subspace of V .

Claim 2.

$$\text{span}(S) = \bigcap_{U \leq V, S \subset U} U$$

Proof. Since $\text{span}(S)$ is a subspace, one inclusion is immediate. For the other, we need to see that if $v \in \text{span}(S)$, then v is in every subspace containing S . This is immediate from the definition of span (linear combinations) and the fact that subspaces are closed under addition and scalar multiplication, and so closed under linear combinations. \square

Infinity.

Exercise 8. If G is an infinite connected graph with every degree finite, then G has an infinite geodesic (an infinite path (v_1, v_2, \dots) such that the distance from v_i to v_j in the graph is $|i - j|$).

Lemma 3.4. (Combinatorial Zorn's Lemma, or König's Lemma.) If T is an infinite tree with all degrees finite, then there is an infinite geodesic ray from the root.

Claim 3. A countable graph is legally c -colorable (has an assignment of colors to vertices so no neighboring vertices share a color) if and only if every finite subgraph is legally c -colorable.

Proof. If the infinite graph is legally c -colorable, of course every finite subgraph is. (Just restrict the colorings.) Conversely: Chooses a nested sequence of subsets of the vertices $= G_0 \subset S_1 \subset S_2 \subset \dots$ so that the union gives all vertices. These give finite subgraphs G_i . Now consider a tree whose n th level consists of colorings of G_i , with an edge between colorings in adjacent levels if the one coloring in level $n + 1$ restricts to the given coloring in level n . The previous lemma gives an infinite geodesic through these colorings, allowing us to color every vertex legally. \square

Given a partial ordered set $(U, <)$, we say that a *chain* is a totally ordered subset, and that a chain, C has an *upper bound* x if $x \geq y$ for every $y \in C$. An element $u \in U$ is called *maximal* if, whenever $v \in U$ is comparable to u , $u \geq v$.

Theorem 3.1. (Zorn's Lemma.) If $(U, <)$ is a partially ordered set, where every chain has an upper bound, then U has a maximal element (an element not smaller than anything in U).

Zorn's Lemma is equivalent to the axiom of choice. We can use it to prove:

Claim 4. Every vector space V has a basis.

Proof. Consider a vector space V , and order all linearly independent subsets by inclusion. Then given a chain of linearly independent subsets, their union is linearly independent and is an upper bound. So there is a maximal element S . As a maximal linearly independent subset, S is basis. \square

For practice using Zorn's lemma, try the following.

Exercise 9. Let $v \neq 0$. Then there is a maximal subspace not containing v .

4. CLASS 4

We'll start today with an application of Zorn's Lemma:

Lemma 4.1. Every graph contains a spanning tree.

Proof. Let G be a graph. We proceed by a Zornification. Let X be the set of subgraphs of G which contain no cycles, ordered by containment. Let $\{U_i \mid i \in I\}$ be some chain of such graphs, and let $\mu = \cup_{i \in I} U_i$. Then μ is an upper bound for the chain.

To see this, note that if e_1, \dots, e_n is a cycle in μ , then for each $1 \leq i \leq n$ we have some $j_i \in I$ for which $e_i \in U_{j_i}$. Because the U_i form a chain, one of the U_{j_i} is maximal amongst them all, so that $U_{j_1} \subset U_{j_i}, U_{j_2} \subset U_{j_i}$ and so on. But then we note that U_{j_i} contains all of the e_i and hence has a cycle. But this contradicts our assumptions. Therefore μ cannot have a cycle.

Now by Zorn's lemma, we know that there is a maximal element T of X . To see that T is a spanning tree, note that it has no cycles, and so could only fail to be a spanning tree if it were not a connected graph. But if that happened, then we could add an edge joining two of the components together without introducing a cycle. This would contradict maximality of T , so we see that T is indeed a spanning tree. \square

Now let's consider the following amazing fact:

Definition 4.1. If K is any field, and n, k are natural numbers, let $M_{n,k}(K)$ be the set of n by k matrices with entries in K . Additionally, let $M_n(K) = M_{n,n}(K)$.

Definition 4.2. If $A \in M_{n,k}(K)$ matrix over any field then the *row-rank* of A is the maximal number of linearly independent rows which A has. Similarly the *column-rank* is the maximal number of linearly independent columns of A .

Theorem 4.1. For any matrix $A \in M_{n,k}(K)$ the row- and column-ranks of M are equal

Proof. We'll come back to this... □

Sign of a Permutation.

Definition 4.3. Let X be a set. Then $\text{Sym}(X)$ denotes the set of permutations of the elements of X . In addition, write $\text{Sym}(1, \dots, n)$ denote the set of permutations on $\{1, \dots, n\}$.

If $p \in \text{Sym}(1, \dots, n)$ then we can consider the pairs (i, j) such that $1 \leq i < j \leq n$; for each such pair we have either $p(i) < p(j)$ or $p(i) > p(j)$.

Definition 4.4. If p is such a permutation, let $\text{inv}(p) = |\{(i, j) \mid 1 \leq i < j \leq n, p(i) > p(j)\}|$, which is the number of pairs whose order is reversed by p .

We say that p is even (odd) if $\text{inv}(p)$ is even (odd).

In addition, define $\text{sign}(p) = (-1)^{\text{inv}(p)}$; so that the sign of p is 1 if $\text{inv}(p)$ is even and -1 if $\text{inv}(p)$ is odd. This is referred to as the *signature* or *sign* of the permutation.

Notice that the only permutation in $\text{Sym}(1, \dots, n)$ which has $\text{inv}(p) = 0$ is the identity permutation. This follows from the fact that if $\text{inv}(p) = 0$ then p is order-preserving so that in particular it sends n to n . This reduces us to the case of a permutation in $\text{Sym}(1, \dots, n-1)$, where we inductively assume the result holds (it is clear in the case $n = 1$, which is the base case).

Theorem 4.2. If $p, q \in \text{Sym}(1, \dots, n)$ then $\text{sign}(pq) = \text{sign}(p) \text{sign}(q)$.

Proof. Define

$$f(x_1, x_2, \dots, x_n) = \prod_{1 \leq i < j \leq n} (x_i - x_j)$$

and note that

$$f(x_{p(1)}, x_{p(2)}, \dots, x_{p(n)}) = \text{sign}(p) f(x_1, x_2, \dots, x_n)$$

because we accumulate a minus sign each time a pair of numbers have their order swapped.

Therefore we have

$$\begin{aligned} \text{sign}(pq) &= \frac{f(x_{p(1)}, x_{pq(2)}, \dots, x_{pq(n)})}{f(x_1, x_2, \dots, x_n)} \\ &= \frac{f(x_{p(1)}, x_{pq(2)}, \dots, x_{pq(n)})}{f(x_{q(1)}, x_{q(2)}, \dots, x_{q(n)})} \frac{f(x_{q(1)}, x_{q(2)}, \dots, x_{q(n)})}{f(x_1, x_2, \dots, x_n)} \\ &= \text{sign}(p) \text{sign}(q) \end{aligned}$$

notice that the left factor is $\text{sign}(p)$ because p permutes the variables $x_{q(1)}, \dots, x_{q(n)}$ to give $x_{pq(1)}, \dots, x_{pq(n)}$; the names of the variables are irrelevant! □

Theorem 4.3. (Bubblesort algorithm) Every permutation in $\text{Sym}(1, \dots, n)$ can be written as a product of transpositions. These may be chosen to be transpositions of adjacent elements.

Proof. Let p be such a permutation. Think of the numbers as labeling bubbles, p as arranging the bubbles by weight. Find i so that $p(i) = n$ (heaviest bubble). Swap i and $i + 1$ (i is a heavier bubble, so i moves downward past $i + 1$). Repeat this process to move the “heaviest bubble” right to the bottom. Now we have $n - 1$ bubbles to sort (or $n - 1$ elements to permute according to p). Inductively, we assume that we can sort these by transposing adjacent elements. \square

By convention, if $M \in M_{n,k}(K)$ we typically denote the element in the i th row and j th column of M by $m_{i,j}$.

Definition 4.5. The *determinant* of $M \in M_n(K)$ is given by

$$\det(M) = \sum_{p \in \text{Sym}(1, \dots, n)} \text{sign}(p) \prod_{i=1}^n m_{i,p(i)}$$

Example 4.

$$\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc$$

because the sign of the identity permutation (corresponding to ad) is 1, while the sign of the permutation $(1, 2)$ (corresponding to bc) is -1 .

Likewise we have

$$\det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} = abc + bfg + cdh - ceg + afh - bdi$$

In computing that last example we made use of the following two facts:

Lemma 4.2. The cyclic permutation (a_1, \dots, a_k) is odd exactly when k is even.

Proof. $(a_1, a_2)(a_1, a_2, \dots, a_k) = (a_1)(a_2, \dots, a_k)$ so that we can decompose the cyclic permutation of k elements as a product of $k - 1$ transpositions. \square

Lemma 4.3. There are equal numbers of even and odd permutations in $\text{Sym}(1, \dots, n)$.

Proof. We can multiply any permutation by $(1, 2)$ to get a permutation of the opposite sign; repeating this will return the original permutation. Therefore this provides a bijection between the odd and even permutations, and there must be the same number of each. \square

The determinant of $M \in M_n(\mathbb{R})$ corresponds to the oriented volume of the shape in n -dimensional space with sides given by the columns of M . In the case $n = 2$ this is a parallelogram, and for $n = 3$ we get a parallelepiped. So, for instance, $ac - bd$ is the oriented area of a parallelogram with vertices $(0, 0)$, (a, c) , (b, c) , $(a + b, c + d)$.

The determinant is zero precisely when the columns of M are linearly dependent. We can think of a parallelepiped as the shape of a cookie; zero volume corresponds to the vectors of the cookie being linearly dependent (i.e. the entire cookie lies in a single plane).

Now we want to see how the determinant of a matrix $M \in M_n(K)$ changes as we change M 's entries.

Lemma 4.4. Multiplying all of the entries in the first row of M by some constant λ will multiply the determinant by λ .

Proof. Each of the rook configurations we can set up in M contains exactly one element from the first row. Therefore the signed products in the determinant arising from each permutation in $\text{Sym}(1, \dots, n)$ will each be multiplied by λ . Therefore their sum (the determinant) is also multiplied by λ . \square

Lemma 4.5. Swapping two rows or columns of M will multiply the determinant by -1 .

Proof. This shifts each of the rook configurations of M to correspond to different permutations by applying a transposition to each, so that the sign of the permutation associated to each product in the determinant changes. \square

It is clear that we can iterate this result to see that permuting the rows or columns of M by a permutation p will multiply the determinant of M by $\text{sign}(p)$.

Lemma 4.6. If two rows of matrix M are equal, then $\det M = 0$.

Proof. In fields of characteristic other than 2, we can swap the two rows to get the same matrix back and note that $\det M = -\det M$ and therefore that $\det M = 0$ from the previous lemma. But for characteristic 2 this deduction is false. We need to note that if the two rows are j th and k th then if p is a permutation we have

$$\prod_{i=1}^n m_{i,p(i)} = \prod_{i=1}^n m_{i,\hat{p}(i)}$$

where (\hat{p}) is the permutation corresponding to the rook configuration which has rooks at $(j, p(k))$ and $(k, p(j))$ instead of $(j, p(j))$ and $k, p(k)$. This permutation is obtained from p by a transposition, so the two have opposite signs and therefore they cancel out in the determinant formula. \square

Lemma 4.7. Adding a row (column) of M to another row (column) of M doesn't change the determinant.

Proof. We may as well assume that we are adding to the first row: we can write

$$M = \begin{bmatrix} v \\ X \end{bmatrix}$$

where v is a row vector and X is a $n - 1$ by n matrix. Then note that we have

$$\det \begin{bmatrix} v + w \\ X \end{bmatrix} = \det \begin{bmatrix} v \\ X \end{bmatrix} + \det \begin{bmatrix} w \\ X \end{bmatrix}$$

by expanding noting that each product in the determinant of the matrix

$$\begin{bmatrix} v + w \\ X \end{bmatrix}$$

has one factor of form $m_{1,j}$ which can be expanded as $v_j + w_j$ where v_j, w_j are the j th entries of v, w . In particular if w is already a row in the matrix then we note that we have

$$\det \begin{bmatrix} w \\ X \end{bmatrix} = 0$$

\square

Similarly we can see that adding λ times one row to another can have no effect on the determinant: by multiplying the j th row by λ , adding the j th row to the i th and then dividing the j th row by λ we obtain a matrix with the same determinant (the first action multiplies the determinant by λ and the last divides it by λ).

This gives us a way to compute determinants efficiently: we can add multiples of rows to each other to kill off entries of the matrix until the computation is easy.

Example 5. To compute the determinant of

$$\begin{bmatrix} 1 & 2 & 0 & 3 \\ -1 & 0 & 0 & 1 \\ 7 & 1 & 8 & 2 \\ 9 & 2 & 3 & 4 \end{bmatrix}$$

we can add row 1 to row 2, -7 times row 1 to row 3, and -9 times row 1 to row 4, and so on.

Remark 1. This process is called *Gaussian elimination*. Ideally, we want to reduce our matrix to something like

$$\begin{bmatrix} a_1 & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & a_2 & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & a_3 & \cdots & \cdots & \cdots \\ \vdots & \vdots & \vdots & \ddots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & a_{n-1} & \cdots \\ 0 & 0 & 0 & \cdots & 0 & a_n \end{bmatrix}$$

in this case the matrix has only got nonzero entries on and above the main diagonal, and is said to be *upper triangular*. In this case the determinant is given by

$$\prod_i a_i$$

because all of the rook configurations other than the diagonal one contain a zero element. This follows from the fact that the rook in the last row must be in the last column to have a nonzero entry, and that the i th row's rook must be in a column greater than the i th, but that all of those other than the i th are needed for the later rows to have nonzero entries.

Lemma 4.8. The number of even and odd permutations are equal (proof in the style of Buddha).

Proof.

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix}$$

All of the products in the determinant of this matrix are ± 1 with sign corresponding to the sign of the permutation. \square

Definition 4.6. If $A \in M_{n,k}(K)$ and $B \in M_{k,m}(K)$ then we define a matrix $AB \in M_{n,m}$ called the *product* of A and B by setting its i, j th entry to be the dot product of the i th row of A and the j th column of B ; i.e denoting this entry by $[AB]_{i,j}$ we have

$$[AB]_{i,j} = \sum_{k=1}^n a_{i,k} b_{k,j}$$

Remark 2. If the dimensions don't agree we have no way to multiply matrices.

Remark 3. The product of two permutation matrices gives a permutation matrix which corresponds to their composite.

Exercise 10. $A(BC) = (AB)C$ holds whenever the product makes sense.

Definition 4.7. The *trace* of a matrix $A \in M_n(K)$ is given by

$$\text{tr}(A) = \sum_{i=1}^n a_{ii}$$

Lemma 4.9. $\text{tr}(AB) = \text{tr}(BA)$.

Proof.

$$\begin{aligned} [AB]_{ii} &= \sum_j a_{ij} b_{ji} \\ \text{tr}(AB) &= \sum_i \sum_j a_{ij} b_{ji} \\ &= \sum_j \sum_i b_{ji} a_{ij} \\ &= \text{tr}(BA) \end{aligned}$$

□

Exercise 11. A matrix $M \in M_n(K)$ has zero determinant if and only if its rows are linearly dependent.

One direction here is straightforward: if a matrix has dependent rows then a nonzero linear combination of them is zero, so one of them is a linear combination of the others. Subtracting that linear combination gives us a zero row but doesn't change the determinant. So the determinant must've been zero all along.

5. CLASS 5

First, a lemma we used when discussing the problem set. This is important in its own right, so I include it here

Lemma 5.1. If $M \in M_n(K)$ has the property that for each n by 1 column vector v we have $Mv = 0$, then M is the zero matrix.

Proof. Let e_i be the column vector whose i th entry is one and whose other entries are all zero. Then Me_i is just the i th column of M . But this is zero by assumption, and so we see that all of M 's columns are zero. □

Definition 5.1. If $A \in M_{n,k}(K)$ then we define the *transpose* of A to be the matrix denoted A^T in $M_{k,n}(K)$ which is specified by setting $[A^T]_{i,j} = [A]_{j,i}$.

Lemma 5.2. For $A \in M_n(K)$, $\det A = \det A^T$.

Proof. First note that the products in the determinant all coincide, so we just need to check the signs of the permutations. However the product corresponding to permutation p in A will correspond to p^{-1} in A^T . Notice that we have

$$\text{sign}(p) \text{sign}(p^{-1}) = \text{sign}(\text{Id}) = 1$$

so that p and p^{-1} have the same sign. □

Theorem 5.1. If $A, B \in M_n(K)$ then $\det(AB) = \det(A) \det(B)$.

Proof. We consider a matrix $M \in M_{2n}(K)$ which is built out of n by n blocks:

$$M = \begin{bmatrix} A & 0 \\ -I & B \end{bmatrix}$$

Notice that $\det(M) = \det(A) \det(B)$ because the only rook configurations which are nonzero are those whose entries in the first n rows are in the first n columns, and therefore that their other entries must be in the last n columns. So the rook configurations which contribute are those given by pairs of rook configurations on n by n matrices.

$$\begin{aligned} \det(M) &= \sum_{p \in \text{Sym}(1, \dots, 2n)} \text{sign}(p) \prod_{i=1}^{2n} m_{i,p(i)} \\ &= \sum_{q, r \in \text{Sym}(1, \dots, n)} \text{sign}(q) \text{sign}(r) \left(\prod_{i=1}^n a_{i,q(i)} \right) \left(\prod_{j=1}^n b_{j,r(j)} \right) \\ &= \sum_{q, r \in \text{Sym}(1, \dots, n)} \left(\text{sign}(q) \prod_{i=1}^n a_{i,q(i)} \right) \left(\text{sign}(r) \prod_{j=1}^n b_{j,r(j)} \right) \\ &= \left(\sum_{q \in \text{Sym}(1, \dots, n)} \text{sign}(q) \prod_{i=1}^n a_{i,q(i)} \right) \left(\sum_{r \in \text{Sym}(1, \dots, n)} \text{sign}(r) \prod_{j=1}^n b_{j,r(j)} \right) \\ &= \det(A) \det(B) \end{aligned}$$

Now add multiples of the first n columns in order to transform our matrix to the form

$$\begin{bmatrix} A & C \\ -I & 0 \end{bmatrix}$$

We get $C = AB$ because to kill off the entry $b_{i,j}$ we need to add $b_{i,j}$ times the i th column of M to the $i + n$ th column of M . Once in this form, by performing n row swaps we can get our matrix to be

$$\begin{bmatrix} -I & 0 \\ A & AB \end{bmatrix}$$

□

which shows our original matrix had determinant

$$(-1)^n \det(-I) \det(AB) = (-1)^{2n} \det(AB) = \det(AB)$$

so that $\det(AB) = \det(A) \det(B)$

Theorem 5.2. Exercise 17 from the problem set:

$$\det \begin{bmatrix} a & b & b & \cdots & b \\ b & a & b & \cdots & b \\ b & b & a & \cdots & b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b & b & b & \cdots & a \end{bmatrix} = (a-b)^{n-1}(a+b(n-1))$$

Proof. Apply column operations: subtract the second from the first, the third from the second, and so on. This yields the matrix

$$\begin{bmatrix} a-b & 0 & 0 & \cdots & b \\ b-a & a-b & 0 & \cdots & b \\ 0 & b-a & a-b & \cdots & b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & a \end{bmatrix}$$

Now add row 1 to row 2, row 2 to row 3, and so on to get

$$\begin{bmatrix} a-b & 0 & 0 & \cdots & b \\ 0 & a-b & 0 & \cdots & 2b \\ 0 & 0 & 0 & \cdots & 3b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & a+(n-1)b \end{bmatrix}$$

which clearly has the desired determinant. \square

In particular this determinant when $a, b > 0$ and $a \neq b$.

Example 6. Suppose we have a town (a set of n people) in which there are some clubs (subsets of the people). Let the clubs be A_1, \dots, A_m , and suppose that $|A_i| = k$ for every i and $|A_i \cap A_j| = l$ for each $i \neq j$.

Theorem 5.3. In the setup above we always have $m \leq n$ (the Fisher inequality).

Proof. Let M be the incidence matrix for the clubs: M 's rows correspond to clubs, and columns correspond to the citizens of the town: i.e. $M_{i,j} = 1$ if person i is in club j , and is 0 otherwise. If $m > n$ then we can add some columns of zeros to get a matrix $A \in M_m(\mathbb{Z})$. Notice then that k is the sum over each row, and that the entry in the (i, j) position of AA^T is k for $i = j$ and l otherwise (since taking the dot product of two rows of A gives us the number of columns for which both of the rows contain a 1). But then AA^T is of the form in the previous theorem $a = k, b = l$ and so the determinant we get is $(k-l)^{n-1}(k+l(n-1))$ which is nonzero unless $k = l$. But in that case we can only have one club. So we note that we have $k \neq l$ and therefore that $\det AA^T \neq 0$. But $\det A = 0$ because we have a zero row, proving $m \leq n$. \square

6. CLASS 6

Lemma 6.1. For matrices consisting of blocks of the form

$$\begin{bmatrix} I & X \\ 0 & I \end{bmatrix}$$

we have

$$\begin{bmatrix} I & X \\ 0 & I \end{bmatrix} \begin{bmatrix} I & Y \\ 0 & I \end{bmatrix} = \begin{bmatrix} I & X+Y \\ 0 & I \end{bmatrix}$$

Definition 6.1. An *elementary matrix* is a matrix in $M_n(K)$ which has 1s on the main diagonal and zeroes in every other entry except one; that is, it is a matrix which differs from the identity matrix in a single off-diagonal entry.

Lemma 6.2. If $A \in M_n(K)$ and E is an elementary matrix then $\det(A) = \det(EA) = \det(AE)$.

Proof. Suppose that E 's nonzero entry is $e_{i,j}$ for $i \neq j$. Then just note that multiplying A on the left by E is the same as adding $e_{i,j}$ times the j th row of A to the i th row. Likewise, multiplying on the right is equivalent to adding $e_{i,j}$ times the i th column of A to the j th column. These operations preserve the determinant. \square

This gives us another way to see why the matrices

$$\begin{bmatrix} A & 0 \\ -I & B \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} A & AB \\ -I & 0 \end{bmatrix}$$

have the same determinant: we can note that we have

$$\begin{bmatrix} A & AB \\ -I & 0 \end{bmatrix} = \begin{bmatrix} A & 0 \\ -I & B \end{bmatrix} \begin{bmatrix} I & B \\ 0 & I \end{bmatrix}$$

and that the matrix

$$X = \begin{bmatrix} I & B \\ 0 & I \end{bmatrix}$$

can be written as a product of elementary matrices by using (Lemma 6.1). But then multiplying by X is equivalent to multiplying by all of the elementary matrices in turn, and each preserves the determinant.

Lemma 6.3. The Vandermonde determinant:

$$\det \begin{bmatrix} 1 & a_1 & a_1^2 & \cdots & a_1^n \\ 1 & a_2 & a_2^2 & \cdots & a_2^n \\ 1 & a_3 & a_3^2 & \cdots & a_3^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & a_n & a_n^2 & \cdots & a_n^n \end{bmatrix} = \prod_{1 \leq j < i \leq n} (a_i - a_j)$$

Proof. Subtract the first row from each of the others. This preserves the determinant and gives

$$\begin{bmatrix} 1 & a_1 & a_1^2 & \cdots & a_1^n \\ 0 & a_2 - a_1 & a_2^2 - a_1^2 & \cdots & a_2^n - a_1^n \\ 0 & a_3 - a_1 & a_3^2 - a_1^2 & \cdots & a_3^n - a_1^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & a_n - a_1 & a_n^2 - a_1^2 & \cdots & a_n^n - a_1^n \end{bmatrix}$$

Now notice that any rook configuration not including the upper left 1 includes some zero entry in the first column, so that we need only consider permutations fixing 1 when we compute the determinant. Thus we have

$$\det \begin{bmatrix} 1 & a_1 & a_1^2 & \cdots & a_1^n \\ 0 & a_2 - a_1 & a_2^2 - a_1^2 & \cdots & a_2^n - a_1^n \\ 0 & a_3 - a_1 & a_3^2 - a_1^2 & \cdots & a_3^n - a_1^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & a_n - a_1 & a_n^2 - a_1^2 & \cdots & a_n^n - a_1^n \end{bmatrix} = \det \begin{bmatrix} a_2 - a_1 & a_2^2 - a_1^2 & \cdots & a_2^n - a_1^n \\ a_3 - a_1 & a_3^2 - a_1^2 & \cdots & a_3^n - a_1^n \\ \vdots & \vdots & \ddots & \vdots \\ a_n - a_1 & a_n^2 - a_1^2 & \cdots & a_n^n - a_1^n \end{bmatrix}$$

For this new matrix the k th row is divisible by $a_k - a_1$ because we have

$$a_k^i - a_1^i = (a_k - a_1)(a_k^{i-1} + a_k^{i-2}a_1 + \cdots + a_k a_1^{i-2} + a_1^{i-1}).$$

Factoring each of the elements of the matrix and pulling the terms out of the determinant, we have

$$\begin{aligned} & \det \begin{bmatrix} a_2 - a_1 & a_2^2 - a_1^2 & \cdots & a_2^n - a_1^n \\ a_3 - a_1 & a_3^2 - a_1^2 & \cdots & a_3^n - a_1^n \\ \vdots & \vdots & \ddots & \vdots \\ a_n - a_1 & a_n^2 - a_1^2 & \cdots & a_n^n - a_1^n \end{bmatrix} \\ &= \prod_{i=2}^n (a_i - a_1) \det \begin{bmatrix} 1 & a_2 + a_1 & \cdots & a_2^{n-1} + a_2^{n-2}a_1 + \cdots + a_2 a_1^{n-2} + a_1^{n-1} \\ 1 & a_3 + a_1 & \cdots & a_3^{n-1} + a_3^{n-2}a_1 + \cdots + a_3 a_1^{n-2} + a_1^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & a_n + a_1 & \cdots & a_n^{n-1} + a_n^{n-2}a_1 + \cdots + a_n a_1^{n-2} + a_1^{n-1} \end{bmatrix} \end{aligned}$$

Now looking at this new matrix we subtract a_1 times the first column from the second, a_2 times the second column from the third, and so on, to get

$$\begin{bmatrix} 1 & a_2 & \cdots & a_2^{n-1} \\ 1 & a_3 & \cdots & a_3^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & a_n & \cdots & a_n^{n-1} \end{bmatrix}.$$

This is simply a Vandermonde matrix of one fewer variables, and so we may proceed by induction. By assumption, the determinant of the smaller matrix is equal to

$$\prod_{2 \leq j < i \leq n} (a_i - a_j),$$

and therefore the original Vandermonde matrix has determinant

$$\left(\prod_{k=2}^n a_k - a_1 \right) \left(\prod_{2 \leq j < i \leq n} (a_i - a_j) \right) = \prod_{1 \leq j < i \leq n} (a_i - a_j)$$

which is the desired result. \square

Definition 6.2. A directed graph is said to be *acyclic* if it has no directed cycles.

Definition 6.3. A *weighted* directed graph is a directed graph G with a function $w: E(G) \rightarrow K$ for some field K ; for us it is safe to assume $K = \mathbb{C}$.

Definition 6.4. If G is a weighted graph, a, b vertices in G , and $p = (e_1, e_2, \dots, e_n)$ a path from a to b then we define the *weight* of p to be

$$w(p) = \prod_{i=1}^n w(e_i).$$

To denote a path p from a to b , we shall sometimes write $p: a \rightarrow b$. We define the *flux* from a to b to be

$$F(a, b) = \sum_{p: a \rightarrow b} w(p)$$

Warning: The following definition of a path configuration is slightly different from the one given in class, in that it now explicitly includes a permutation; this does not change the content of the theorems at all, but it makes the definition easier to write out.

Definition 6.5. Given G , a finite acyclic weighted directed graph with weight function w , and a_1, \dots, a_n and b_1, \dots, b_n subsets of the vertices, then a *path configuration* P on G with respect to these vertices is a set p_1, \dots, p_n of n paths in G together with a permutation $q \in \text{Sym}(1, \dots, n)$ such that for each i , p_i is a path from a_i to $b_{q(i)}$.

We define the *sign* of P to be the sign of q .

We call a path system P *vertex-disjoint* if each vertex in the graph is passed through by at most one of the paths in P .

Theorem 6.1. If G is a finite acyclic weighted directed graph with weight function w , choose vertices a_1, \dots, a_n and b_1, \dots, b_n to be sets of sources and sinks, respectively. Define a matrix M by $m_{i,j} = F(a_i, b_j)$. Then

$$\det M = \sum_{\text{path systems } P} \text{sign}(P) \prod_{i=1}^n W(p_i) = \sum_{\substack{\text{vertex-disjoint} \\ \text{path systems } P}} \text{sign}(P) \prod_{i=1}^n W(p_i)$$

Proof. Fix some permutation q and look at what happens with the rook configuration corresponding to that permutation: we have

$$\begin{aligned} \prod_{i=1}^n F(a_i, b_{q(i)}) &= \prod_{i=1}^n \left(\sum_{\substack{\text{paths } p_i \\ a_i \rightarrow b_{q(i)}}} w(p_i) \right) \\ &= \sum_{\substack{(p_1, p_2, \dots, p_n) \\ p_i \text{ paths } a_i \rightarrow b_{q(i)}}} \prod_{i=1}^n w(p_i) \end{aligned}$$

But then summing over all of the rook configurations and taking into account permutation signs we get

$$\det M = \sum_{\text{path systems } P} \text{sign}(P) \prod_{i=1}^n W(p_i).$$

Now it remains to see that we can ignore non-vertex-disjoint path systems. To see this, take P to be such a path system. Define another path system \hat{P} as follows: find the least i so that p_i crosses another path. Take the least j for which p_i intersects p_j ; define new paths \hat{p}_i and \hat{p}_j by swapping p_i and p_j over at the first point where they meet; i.e if p_i is (u_1, \dots, u_m) , p_j is (v_1, \dots, v_n) , and the first time they intersect is when v_k and u_l both go into the same vertex, then we set

$$\begin{aligned} \hat{p}_i &= (u_1, \dots, u_k, v_{l+1}, \dots, u_n) \\ \hat{p}_j &= (v_1, \dots, v_l, u_{k+1}, \dots, v_m) \end{aligned}$$

for $k \neq i, j$ set $\hat{p}_k = p_k$. Let $\hat{P} = (\hat{p}_1, \dots, \hat{p}_n)$. We have $w(P) = w(\hat{P})$, $\text{sign}(P) = -\text{sign}(\hat{P})$, and $(\hat{\hat{P}}) = P$, so in the formula, each non-vertex-disjoint path system has a corresponding system which exactly cancels out its contribution. Thus, we need only consider vertex-disjoint systems. \square

Theorem 6.2. The Cauchy-Binet formula: if $n < k$, $A \in M_{n,k}(K)$, $B \in M_{k,n}(K)$ then we have

$$\det(AB) = \sum_{\substack{S \subseteq \{1, \dots, k\} \\ |S|=n}} \det(A_S) \det(B_S)$$

where A_S is the n by n submatrix of A given by the columns with indices in S , B_S is the n by n submatrix of B given by the rows with indices in S .

Proof. Let G be the directed graph with vertices $a_1, \dots, a_n, b_1, \dots, b_k, c_1, \dots, c_n$ and edges $e_{i,l}$ from each a_i to each c_l , and $f_{l,j}$ from each c_l to each b_j , with $w(e_{i,l}) = a_{i,l}$, $w(f_{l,j}) = b_{l,j}$. Then applying the path configuration theorem to this graph. Fix a set S of n vertices from the c_l ; notice that such sets correspond to a set of vertex-disjoint path configurations P (all

vertex-disjoint P pass through precisely n of the c_l). Then

$$\det(A_S) \det(B_S) = \sum_{\substack{\text{vertex disjoint } P \\ \text{passing through } S}} \text{sign}(P) \prod_{i=1}^n w(p_i)$$

and summing over all of the S gives the result. \square

Definition 6.6. If G is an undirected graph with no loops, define an incidence matrix of $I(G)$ for G to be a matrix whose i th row is labeled by the vertex v_i of G and whose columns are labeled by the edges of G , such that each column labeled by e contains only two nonzero entries which are 1 and -1 , and these entries are in the rows corresponding to the vertices e joins.

Example 7.

$$\begin{bmatrix} 1 & 0 & -1 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix}$$

is an incidence matrix for the graph which is a “triangle”.

Notice that $L = I(G)I(G)^T$ is a matrix with

$$l_{i,i} = \deg(v_i)$$

and

$$l_{i,j} = -|\{\text{edges } e \text{ joining } v_i \text{ to } v_j\}|$$

Theorem 6.3. If G is an undirected simple graph then if $L = I(G)I(G)^T$, and $L_{1,1}$ denotes the matrix obtained by deleting the first row and column of L , then $\det(L_{1,1})$ is the number of spanning trees in G . This number is n^{n-2} .

Proof. Not yet! \square

7. CLASS 7

Review. Last time: Cauchy-Binet, which says if A is $m \times n$ and B is $n \times m$, then:

$$\det(AB) = \sum_S \det(A_S) \det(B_S),$$

where the sum is over all S subsets with m elements of $\{1, \dots, n\}$.

Recall:

Definition 7.1. If G is a directed graph with no loops, define an incidence matrix of $I(G)$ for G to be a matrix whose i th row is labeled by the vertex v_i of G and whose columns are labeled by the edges of G , such that each column labeled by e contains only two nonzero entries which are 1 and -1 , and these entries are in the rows corresponding to the vertices e joins.

Notice that $L = I(G)I(G)^T$ is the degree matrix minus the adjacency matrix, and we call it L , the Laplacian. Note that $\det(L)$ is always 0 (summing all columns gives the zero vector), so is not an interesting number. A more interesting number:

Recall the notion that $L_{1,1}$ is the matrix obtained by deleting the first row and first column from G .

Matrix Tree Theorem.

Theorem 7.1. The number of different spanning trees of a connected graph G is $\det(L_{1,1})$.

Proof. Suppose that G has n vertices and m edges, and let J be the incidence matrix with the first row removed. Then $L_{1,1}$ is JJ^T . We will use the Cauchy-Binet formula, of course, which tells us

$$\det(L_{1,1}) = \sum \det(J_S) \det(J_S^T),$$

where $S \subset \{1, \dots, m\}$ has size $n - 1$, specifying which columns (representing edges) are kept. We need to show that $\det(J_S)$ is ± 1 if the edges in S form a spanning tree, and is 0 otherwise.

Suppose these $n - 1$ edges do not form a spanning tree. Then the graph G' consisting of G 's n vertices with these $n - 1$ edges is disconnected, so has at least 2 components. One of these components does not contain the vertex v_1 (corresponding to the first row we removed). When we sum the rows corresponding to this component, we get zero.

Suppose the subgraph G' is a tree. We can permute the rows and columns so that the resulting matrix has 1's on the diagonal, and is lower triangular. This shows that the determinant is ± 1 . We can obtain these permutations by choosing an appropriate order on the vertices and edges: Choose a leaf (call it v_1) and its corresponding edge (call it e_1). Now remove these. Choose a leaf in the new graph, call it v_2 , and so on.

□

Application to K_n . For K_n ,

$$L = \begin{bmatrix} n-1 & -1 & -1 & \cdots & -1 \\ -1 & n-1 & -1 & \cdots & -1 \\ -1 & -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & -1 & \cdots & n-1 \end{bmatrix}.$$

From the exercise, the determinant is n^{n-2} .

More Linear Algebra. Let V, W be vector spaces over a field k .

Definition 7.2. A map $\phi : V \rightarrow W$ is *linear* if $\phi(\alpha v_1) = \alpha \phi(v_1)$, and $\phi(v_1 + v_2) = \phi(v_1) + \phi(v_2)$ for all $\alpha \in k$ and all $v_1, v_2 \in V$.

If $\phi : V \rightarrow W$ is a linear map, we can also define two important sets associated to ϕ .

Definition 7.3. The *kernel* of ϕ is the set $\ker(\phi) = \{v \in V \mid \phi(v) = 0\}$.

Definition 7.4. The *image* of ϕ is the set $\text{im}(\phi) = \phi(V) = \{w \in W \mid \exists v \in V \text{ s.t. } \phi(v) = w\}$.

The image and kernel are subspaces.

Lemma 7.1. If S spans V , then $\phi(S)$ spans $\phi(V)$.

Lemma 7.2. Let B be a basis for V , and let $f : B \rightarrow W$. Then there is a unique linear map $\phi : V \rightarrow W$ which extends f .

Lemma 7.3. Let $\phi : V \rightarrow W$ be linear. Then the following are equivalent:

- (1) ϕ is injective

- (2) $\ker(\phi) = 0$
- (3) $S \subset V$ is linearly independent implies $\phi(S)$ is.

Lemma 7.4. $\dim(V) = \dim(W)$ if and only if V is isomorphic to W .

8. CLASS 8

Question 1. How is every matrix a linear map? If A is a $n \times k$ matrix, then A determines a function $f_A: K^k \rightarrow K^n$ given by $f_A(v) = A \cdot v \in K^n$, where $v \in K^k$.

Definition 8.1. A *linear transformation* is a linear map $V \rightarrow V$. (Note that the source and target are the same.)

Example 8. (1) Euclidean transformations are linear transformations.

- (2) Fix a basis, and ‘stretch’ (by a possibly different amount) along each basis vector. (These are the *diagonalizable* linear transformations.)
- (3) Rotations (most of these are *not* diagonalizable).

Question 2. What does a linear transformation look like? We’ll try to decompose the linear transformation into ‘indecomposable’ parts, and then understand the indecomposable linear transformations.

Definition 8.2. Let $\varphi: V \rightarrow V$ be a linear transformations, and let U be a subspace of V . Then U is *invariant* (or φ -*invariant*) if for all $u \in U$, $\varphi(u) \in U$.

Example 9. If $\varphi: V \rightarrow V$ is a rotation along some axis, what are the invariant subspaces of φ ? The axis of rotation and the plane perpendicular to the axis of rotation.

Definition 8.3. If $\varphi: V \rightarrow V$ is a linear transformation, then an *eigenvector* of φ is a non-zero vector v such that $\varphi(v) = \lambda v$ for some $\lambda \in K$. λ is known as the *eigenvalue* of v . (Note that this is equivalent to saying that the span of v is φ -invariant. So, eigenvectors are the same thing as one-dimensional invariant subspaces.) Note that we allow $\lambda = 0$.

Note that if U is an invariant subspace of φ , then φ restricts to a linear transformation from U to itself. Now, the question arises: if we’ve found one invariant subspace, what happens to the rest of the space? For this, quotient spaces would be nice; we’ll come back to it later.

Today, we’ll look at symmetric matrices, and we’ll see that they’re ‘nicer’ than arbitrary matrices. We’ll define the standard inner product on \mathbb{R}^n ; later, we’ll introduce a generalization, just as vector spaces are generalizations of \mathbb{R}^n .

Definition 8.4. If $v, w \in K^n$, then the *standard inner product* of v and w is

$$\langle v, w \rangle = v^t \cdot w \in K.$$

(Here, we’re thinking of v and w as $n \times 1$ column vectors).

Now, we’ll express some geometric notions in terms of the inner product.

- (1) What does it mean if $\langle v, w \rangle = 0$? This happens if and only if v and w are orthogonal (denoted $v \perp w$).
- (2) What is $\langle v, v \rangle$? It’s the square of the Euclidean length of v . So, v has length 1 if and only if $\langle v, v \rangle = 1$.

Definition 8.5. We say that a basis $\{b_1, \dots, b_n\}$ of V is *orthogonal* if $\langle b_i, b_j \rangle = 0$ for $i \neq j$, and we say that an orthogonal basis is *orthonormal* if, in addition, $\langle b_i, b_i \rangle = 1$. (Note that the orthonormal bases of \mathbb{R}^n are just rotations of the standard basis.)

The goal for the day is the following theorem.

Theorem 8.1 (Spectral Theorem). If $A = A^t$ is an $n \times n$ real matrix, then there exists an orthonormal basis $\{b_1, \dots, b_n\}$ and constants $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that $Ab_i = \lambda_i b_i$ for all i . (In other words, every symmetric matrix admits an orthonormal basis of eigenvectors.)

Lemma 8.1. For any matrices A and B (where the product makes sense),

$$(A \cdot B)^t = B^t \cdot A^t.$$

Proof. This is true by definition (just check by direct computation). \square

Definition 8.6. If $S \subset V$ is any subset, then we define the *orthogonal complement* of S to be the set

$$S^\perp = \{v \in V \mid v \perp s \text{ for all } s \in S\}.$$

Note that even if S is not a subspaces of V , S^\perp is (because the standard inner product is *bilinear*).

Lemma 8.2. Every orthonormal system is linearly independent.

Proof. Suppose that $\{u_1, \dots, u_k\}$ is orthonormal, and assume that $\sum \alpha_i u_i = 0$. Then,

$$0 = \left\langle \sum \alpha_i u_i, u_j \right\rangle = \alpha_j.$$

\square

Lemma 8.3. If U is a subspace of \mathbb{R}^n , then U has an orthonormal basis.

Proof. We'll inductively construct increasingly large orthonormal subsets of U . By the previous lemma, they're all linearly independent, so eventually we'll get to a basis. Clearly we can choose an orthonormal subset of U of size 1; just choose any vector of norm 1. Now, assume we've constructed an orthonormal system $\{u_1, \dots, u_k\}$ of U . If it doesn't span the whole space, let v be in U , but outside the span of $\{u_1, \dots, u_k\}$. Then, let $\lambda_i = \langle v, u_i \rangle$, and let $v_0 = \sum \lambda_i u_i$. Note that $v - v_0$ is also outside the span of $\{u_1, \dots, u_k\}$. Moreover, for all $u_i \in U$,

$$\langle v - v_0, u_i \rangle = \langle v, u_i \rangle - \langle v_0, u_i \rangle = \lambda_i - \left\langle \sum_{j=1}^k \lambda_j u_j, u_i \right\rangle = \lambda_i - \sum_{j=1}^k \lambda_j \langle u_j, u_i \rangle = \lambda_i - \lambda_i = 0.$$

So, $v - v_0$ is orthogonal to every u_i . Thus, we can divide $v - v_0$ by its norm to obtain a vector u_{k+1} of norm 1 (still orthogonal to every u_i , $1 \leq i \leq k$), and we obtain a larger orthonormal set $\{u_1, \dots, u_k, u_{k+1}\}$. \square

Lemma 8.4. If U is a subspace of V , then

- (1) $\dim U^\perp + \dim U = \dim V$, and
- (2) $(U^\perp)^\perp = U$.

Proof. For the first part, first note that $U \cap U^\perp = \{0\}$, since if $v \in U \cap U^\perp$, then $\langle v, v \rangle = 0$, so $v = 0$. Thus, $\dim U + \dim U^\perp \leq \dim V$. We know that

$$\dim V \geq \dim(U + U^\perp) = \dim U + \dim U^\perp - \dim(U \cap U^\perp) = \dim U + \dim U^\perp,$$

so it suffices to show that

$$U + U^\perp = \{v + v' \mid v \in U, v' \in U^\perp\} = V.$$

To see this, let $v \in V$. Let $\{u_1, \dots, u_k\}$ be an orthonormal basis of U , and let $\lambda_i = \langle v, u_i \rangle$, let $v_0 = \sum \lambda_i u_i$. Then,

$$v = v_0 + (v - v_0),$$

and $v_0 \in U$, $v - v_0 \in U^\perp$, since for all $u_i \in U$,

$$\langle v - v_0, u_i \rangle = \langle v, u_i \rangle - \langle v_0, u_i \rangle = \lambda_i - \left\langle \sum_{j=1}^k \lambda_j u_j, u_i \right\rangle = \lambda_i - \sum_{j=1}^k \lambda_j \langle u_j, u_i \rangle = \lambda_i - \lambda_i = 0.$$

So, $v \in U + U^\perp$, as desired.

Now we'll do the second part. Clearly $U \subset (U^\perp)^\perp$, since if $u \in U$, for all $w \in U^\perp$, $\langle v, w \rangle = 0$. On the other hand, $\dim U^\perp = \dim V - \dim U$, so

$$\dim(U^\perp)^\perp = \dim V - \dim U^\perp = \dim V - (\dim V - \dim U) = \dim U,$$

so $U = (U^\perp)^\perp$ □

Corollary 8.1.1. If S is any subset of V , then $(S^\perp)^\perp$ is the span of S .

Question 3. Does every vector space have an orthonormal basis? Well, as stated this question doesn't really make sense, since not every vector space comes with an inner product. So, we're really asking if every inner product space has an orthonormal basis. (Note that we're talking about infinite-dimensional spaces here, since we answered the question above for finite dimensional spaces.) What properties should our inner product have? It should be bilinear, positive definite, and symmetric. We have to use Zorn's lemma, since our space may not have a countable basis. We use Zorn's lemma on orthonormal systems with respect to inclusion. Every chain has an upper bound, by taking the union. So, we have maximal elements. Why does it span the space? We're not really sure, because if we don't span, we may not be able add another *orthogonal* vector.

From here on out, we're in the setting of the Spectral Theorem, so notation and assumptions are as in the statement of that theorem.

Lemma 8.5. If U is A -invariant, then U^\perp is A^t -invariant.

Proof. Let $v \in U^\perp$, and let $u \in U$. Then,

$$\langle u, A^t v \rangle = u^t A^t v = (Au)^t v = 0,$$

since $Au \in U$ and $v \in U^\perp$.

Since A is a symmetric matrix in the setting of the spectral theorem, then $A = A^t$. So, if U is A -invariant, then U^\perp is A -invariant as well. □

Lemma 8.6 ('Ugly'). Define

$$f(t) = \frac{a + bt + ct^2}{d + et^2},$$

where $d > 0$. If 0 is a maximum point of f , i.e., if $f(0)$ is a maximum, then $b = 0$.

Proof. The ‘ugly’ proof is: take the derivative, you know that it must be zero at zero; this implies $b = 0$.

The conceptual proof is: near zero, the t^2 parts don’t matter, so f basically looks like a line; it can only have a maximum at zero if it has slope zero. \square

Proof. (Of the Spectral Theorem).

We claim that for any A -invariant subspace U , there exists a non-zero eigenvector $u \in U$ of A . If we prove this claim, we’ll be done: pick an eigenvector from V ; look at the orthogonal complement of this eigenvector, pick an eigenvector from here; take the orthogonal complement of the first two eigenvectors ... until we can’t any more. This will give an orthonormal basis of V consisting of eigenvectors for A .

Define the Rayleigh-quotient:

$$R(x) = \frac{\langle x, Ax \rangle}{\langle x, x \rangle} = \frac{x^t Ax}{x^t x},$$

where $x \in U$ is non-zero. Note that $R(\lambda x) = R(x)$ for all $\lambda \in \mathbb{R}^\times$. So, if we know the value of R on the unit sphere, we know the value everywhere. Let B be the unit sphere, i.e.,

$$B = \{x \in U \mid \langle x, x \rangle = x^t x = 1\}.$$

Now, B is compact (being closed and bounded) and R is continuous, so there exists $u \in B$ such that $R(x) \leq R(u)$ for all $x \in B$. We claim that u is an eigenvector. To do this, we’ll show that u^\perp is A -invariant. Then this means that the space spanned by u is also A -invariant, which means exactly that u is an eigenvector of A . So, let $v \perp u$, and define

$$\begin{aligned} f(t) &= R(u + tv) = \frac{\langle (u + tv), A(u + tv) \rangle}{\langle u + tv, u + tv \rangle} \\ &= \frac{\langle u, Au \rangle + \langle v, Au \rangle t + \langle u, Av \rangle t + \langle v, Av \rangle t^2}{\langle u, u \rangle + \langle u, v \rangle t + \langle v, u \rangle t + \langle v, v \rangle t^2} \\ &= \frac{\langle u, Au \rangle + 2\langle u, Av \rangle t + \langle v, Av \rangle t^2}{\langle u, u \rangle + \langle v, v \rangle t^2}. \end{aligned}$$

(Note that in the above manipulations, we used that $\langle v, Au \rangle = \langle u, Av \rangle$, which is true because A is symmetric.)

Now, we know that f has a maximum at zero, since R has a maximum at u . Moreover, $\langle u, u \rangle > 0$. Thus, by the ‘ugly’ lemma, $2\langle u, Av \rangle = 0$, i.e., $\langle u, Av \rangle = 0$, and u^\perp is A -invariant, as claimed. \square

Exercise 12. Let G be an undirected graph on n points with adjacency matrix $A = A(G)$. (So $A = A^t$.) By the Spectral Theorem, A admits a basis of eigenvectors; let $\lambda_0(G) \geq \lambda_1(G) \geq \dots \geq \lambda_{n-1}(G)$ be the eigenvalues of A . (We call these the eigenvalues of G .)

Oh no! Are “the eigenvalues of a symmetric matrix ” well-defined? We only showed that there is an orthonormal eigenbasis; how do we know that someone else can’t find a different eigenbasis that admits different eigenvalues? One approach is to show that the eigenvalues are precisely the roots of the characteristic polynomial, which is independent of any choice of basis. What’s another approach? Well, the eigenvectors of eigenvalue λ are precisely the elements of the kernel of $A - \lambda I$. We call these the *eigenspaces* of A , and we denote it V_λ . In other words,

$$V_\lambda = \{v \in V \mid Av = \lambda v\}.$$

If $\lambda \neq \lambda'$, then $V_\lambda \cap \ker V_{\lambda'} = \{0\}$. Moreover, if $\lambda \neq \lambda'$, then $V_\lambda \perp V_{\lambda'}$: if $v \in V_\lambda$ and $v' \in V_{\lambda'}$, then

$$\lambda \langle v, v' \rangle = \langle \lambda v, v' \rangle = \langle Av, v' \rangle = \langle v, Av' \rangle = \langle v, \lambda' v' \rangle = \lambda' \langle v, v' \rangle;$$

since $\lambda \neq \lambda'$, we must have $\langle v, v' \rangle = 0$. So, if $\lambda_1, \dots, \lambda_j$ are the eigenvalues of A with respect to some basis, then

$$V = V_{\lambda_1} + \dots + V_{\lambda_j}.$$

Now, if μ is some other ‘impostor’ eigenvalue, then we should have V_μ is perpendicular to every V_{λ_i} ; but the sum of these already spans V , so V_μ must be zero. Thus, the eigenvalues are uniquely determined by the linear transformation.

9. CLASS 9

Some linear algebra.

Question 4. If $A, B \in M_n(\mathbb{R})$ and $AB = I$, does it follow that $BA = I$?

If we think of the matrices as linear transformations from \mathbb{R}^n to \mathbb{R}^n , then if $B(v) = w$, we have $A(w) = v$, so $BA(w) = B(v) = w$, and hence BA is the identity when we restrict to the image of B . However, since $AB = I$, we must have that B is injective. Since injective linear maps send bases to linearly independent sets, we have that the image of B is at least n dimensional, and hence B is surjective.

If we have an orthonormal basis of a vector space, $\{u_1, \dots, u_n\}$ and we let $U = [u_1 | \dots | u_n]$ be the matrix whose columns are the u_i , then the condition that we have an orthonormal basis is equivalent to $U^T U = I$, which by the above observation is equivalent to $U U^T = I$, so, the rows of U also give us an orthonormal basis.

Eigenvalues of the adjacency matrix. Let G be an undirected graph on n points, A the adjacency matrix of G , and let $\lambda_0(G) \geq \lambda_1(G) \geq \dots \geq \lambda_{n-1}(G)$ be the eigenvalues of A . Multiplying a vector (whose entries represent quantities stored at each vertex) by the adjacency matrix distributes the quantities around by sending the amount at vertex v_i to each of its neighbors while simultaneously taking in the sum of the quantities from all neighbors. We can think of the situation like that of a town of forgetful gossips: they tell their neighbors all of their secrets, but only remember what they have just been told themselves.

Exercise 13. If G is d -regular, then $\lambda_0(G) = d$.

Proof. We must show both that d is an eigenvalue and that no eigenvalue is larger. We have that

$$(9.1) \quad A \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} d \\ \vdots \\ d \end{pmatrix}$$

because multiplying this particular vector by the adjacency matrix gives the number of incoming neighbors. If we have some other nonzero eigenvector $v = (v_1, \dots, v_n)^T$ and let i be the index of the vertex of largest absolute value, so $|v_i| \geq |v_j|$ for $i \neq j$. Then

$$|(Av)_i| = \left| \sum_{j \rightarrow i} v_j \right| \leq \sum_{j \rightarrow i} |v_j| \leq d |v_i|.$$

Thus, the corresponding eigenvalue is at most d . □

Note that we have proved even more than the statement of the problem: If G is d -regular, then no eigenvalue has absolute value larger than d .

Lemma 9.1. If G is d -regular, then $\lambda_1(G) = d$ if and only if G is not connected.

Proof. Suppose that $v = (v_1, \dots, v_n)^T$ is a nonzero eigenvector with eigenvalue d . Let i be an index such that $v_i \geq v_j$ for $i \neq j$. By considering $-v$ if necessary, we may assume that $v_i > 0$. Then since $0 = (dv - Av)_i = dv_i - \sum_k v_{j_k} = \sum_k (v_i - v_{j_k})$ where the j_k are the vertices connected by an edge to i . By the maximality of v_i , we have that all the terms in the sum are non-positive, and thus are zero. Thus, $v_{j_k} = v_i$ for all k . If G is connected, then this argument shows that $v_i = v_j$ for every j , so the eigenspace is one dimensional.

Conversely, if G is disconnected, then pick a component and let $v_i = 1$ for i in the component, and 0 otherwise. This shows that the dimension of V_d is at least the number of components (and the argument above shows that the dimension is exactly the number of components). □

This suggests that the *eigenvalue gap* $\lambda_0(G) - \lambda_1(G)$ gives some kind of indication of the connectedness of a graph. Another useful invariant which gives information about the connectedness of the graph is the *spectral radius*, $\rho(G) = \max_{0 < i < n} |\lambda_i(G)|$. If G is d -regular, then we have that

$$d \geq \rho(G) \geq 2\sqrt{d-1} = \rho(T_d)$$

where T_d is the d -regular tree.

Theorem 9.1 (Alon-Boppana). If G_n is a sequence of d -regular graphs and $|G_n|$ tends towards infinity, then $\liminf \rho(G_n) \geq 2\sqrt{d-1}$.

It is an open problem whether for any given d , there are infinitely many G with $\rho(G) < 2\sqrt{d-1}$.

Theorem 9.2. $\text{diam}(G) \leq \frac{\log(n-1)}{\log(d/\rho(G))}$ if G is d regular on n points.

How can we prove this? Naively, we know that there are at most $d(d-1)^k$ vertices of distance k away from a given vertex, so we expect a logarithmic bound on diameter in the best case scenario. This result shows that this is indeed the case when the spectral radius is small. We will need a preliminary result before we can prove the theorem.

Theorem 9.3 (Cauchy-Schwarz). If $x, y \in \mathbb{R}^n$, then $\langle x, y \rangle^2 \leq \langle x, x \rangle \langle y, y \rangle$.

Proof. Since the equation is trivial if either x or y is zero, let us assume otherwise. Let Λ be an unspecified parameter. Then

$$0 \leq \langle x - \Lambda y, x - \Lambda y \rangle = \langle x, x \rangle + \Lambda^2 \langle y, y \rangle - 2\Lambda \langle x, y \rangle$$

If we let $\Lambda = \sqrt{\frac{\langle x, x \rangle}{\langle y, y \rangle}}$ and divide the above equation by Λ we get

$$2\langle x, y \rangle \leq \langle x, x \rangle / \Lambda + \Lambda \langle y, y \rangle = 2\sqrt{\langle x, x \rangle \langle y, y \rangle}.$$

This gives the desired inequality. □

Remark 4. Note that, by the theorem of arithmetic and geometric means (or calculus), our choice of Λ gives the smallest possible value for the right hand side, so no better result is possible by picking a different value of Λ .

Now, we can prove the theorem.

Proof. Let A be the adjacency matrix of G , and let u_0, \dots, u_{n-1} be an orthonormal basis of eigenvectors with eigenvalues $\lambda_i(G)$, so that $Au_i = \lambda_i u_i$, $u_i^T u_i = 1$, and $u_i^T u_j = 0$ when $i \neq j$. Let $U_i = u_i u_i^T$. These U_i have several interesting properties:

- $U_i U_j = u_i u_i^T u_j u_j^T = u_i 0 u_j^T = 0$
- $U_i U_i = u_i u_i^T u_i u_i^T = u_i u_i^T = U_i$
- $AU_i = A u_i u_i^T = \lambda_i U_i$
- $\sum \lambda_i U_i = A$
- $\sum U_i = I$

We can see the last two properties by applying both sides to each of the u_j and noting that, if two linear maps agree on a basis, then they are the same.

Additionally, we note that $(A^n)_{ij}$ is the number of n -step paths from i to j .

Therefore, if all the entries of A^m are nonzero, then the $\text{diam}(G) \leq m$. However, using the properties of the U_i we have

$$A^m = \left(\sum \lambda_i U_i \right)^m = \sum \lambda_i^m U_i$$

Because $(1, 1, \dots, 1)^T$ has eigenvalue $\lambda_0(G)$, we can choose $u_0 = (1/\sqrt{n}, 1/\sqrt{n}, \dots, 1/\sqrt{n})^T$, and hence

$$U_0 = \begin{pmatrix} 1/n & \cdots & 1/n \\ \vdots & \ddots & \vdots \\ 1/n & \cdots & 1/n \end{pmatrix}$$

Using this and the fact that $\lambda_0 = d$, as well as the definition of the spectral radius ρ , we have

$$\begin{aligned} |(A^m)_{rs}| &= \left| d^m/n + \sum_{i>0} \lambda_i^m (U_i)_{rs} \right| \\ &\geq d^m/n - \left| \sum_{i>0} \lambda_i^m (U_i)_{rs} \right| \\ &= d^m/n - \rho^m \left| \sum_{i>0} (\lambda_i/\rho)^m (U_i)_{rs} \right| \\ &\geq d^m/n - \rho^m \sum_{i>0} |(U_i)_{rs}| \end{aligned}$$

Let us estimate $\sum_{i>0} |(U_i)_{rs}|$. Taking the sum to include $i = 0$ as well just increases the value by $1/n$, so we will estimate $\sum_i |(U_i)_{rs}|$. If we let u_{ij} denote the j th coordinate of u_i , then by Cauchy-Schwarz, $\sum_i |(U_i)_{rs}| = \sum_i |u_{ir} u_{is}| \leq \sqrt{(\sum_i u_{ir}^2)(\sum_i u_{is}^2)}$. However, if we let U be orthogonal the matrix whose columns are the u_i , then the j th row is (u_{1j}, \dots, u_{nj}) , and since the rows of U form an orthonormal basis, $\sqrt{(\sum_i u_{ir}^2)(\sum_i u_{is}^2)} = 1$.

Therefore, $(A^m)_{rs} \geq d^m/n - \rho^m(1 - 1/n)$ which is greater than zero if $d^m/n > \rho^m(1 - 1/n) \Leftrightarrow (d/\rho)^m > n - 1 \Leftrightarrow m \geq \log(n - 1)/\log(d/\rho)$ \square

The theorem is useless if $\rho = d$. There are two ways that this can happen. The first, as we have seen, is if $\lambda_1 = d$, which happens if G is not connected. The second is if $(-d)$ is an eigenvalue.

Exercise 14. If G is d regular, then the adjacency matrix has an eigenvalue of $(-d)$ if and only if G is bipartite.

Of course, if G is bipartite, then A^m will have zeros in it for any m , because after an even number of steps, one ends up in the part one started in, and in an odd number of steps, one ends up in the other part. Therefore, we see that in the two cases that the theorem tells us nothing, the method of proof has no chance of yielding interesting information.