Ghostly graph theory

Travis, Mathcamp 2024

For this entire class, G is a connected graph with n vertices.

1. adjacency matrices and walks

1.1. a humble beginning

Definition. A *walk of length* k is a sequence of vertices $x_0, x_1, \ldots, x_k$ in G such that x_{i+1} is adjacent to x_i. (The vertices do not have to be distinct.) A walk is called *closed* if $x_k = x_0$.

Problem 1. How many walks of length k does K_4 have? What about K_n?

Here is a harder question: How many *closed* walks of length k does K_n have? Or more generally:

Question. Given a graph G, how can we find the number of closed walks of length k? (For every k.)

This is the question we'll answer today, and we'll use *linear algebra* to do it.

Definition. The *adjacency matrix* of G is the $n \times n$ matrix A whose entries are

$$A_{i,j} = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are connected} \\ 0 & \text{if they are not.} \end{cases}$$

Here is an example graph and its adjacency matrix:

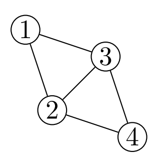

$$\begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}$$

The adjacency matrix is just another way to think about the information a graph contains, but a table of numbers is harder to understand than a picture. But the whole point of this class is that thinking of it as a *matrix*, not just a table of numbers, will be super, super useful.

Problem 2.

(a) Prove that $(A^k)_{i,j}$, the (i,j) entry of $\underbrace{A \cdot A \cdots A}_{k}$, is the number of walks of length k that start at i and end at j.

(b) Prove that $\operatorname{tr}(A^k)$ is the number of closed walks of length k in G. (Remember that $\operatorname{tr}(M)$ is the sum of the diagonal entries of M.)

Aha! So we need to think about the trace of a matrix. Let's collect some linear algebra.

1.2. gathering some tools

For this entire class, M is an $n \times n$ matrix.

Definition. A vector $v \neq \mathbf{0}$ is an *eigenvector* of M if there is a real number λ such that $Mv = \lambda \cdot v$, in which case λ is called the *eigenvalue* associated with M.

It is possible that some of the eigenvalues are complex numbers. (For example, the matrix $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ has complex eigenvalues.)

Problem 3. (Basic properties of eigenvectors)

(a) If v_1 and v_2 are eigenvectors with eigenvalue λ, then $av_1 + bv_2$ is an eigenvector with eigenvalue λ, for any $a, b \in \mathbb{R}$.

(b) If $v_1, v_2, \ldots, v_k$ are eigenvectors and each has a different eigenvalue, then these vectors are linearly independent. [Hint: Suppose there were a linear dependence, and apply A to it.]

(c) M has at most n different eigenvalues.

This is (part 1 of) the Big Theorem we will reference throughout the entire class.

Theorem (Spectral theorem, part 1). *If M is a symmetric matrix, then there is a basis $v_1, \ldots, v_n$ of $\mathbb{R}^n$ where each v_i is an eigenvector of M. (This is called an* eigenbasis *of M.) Every eigenvalue associated to these vectors is a* ***real number****, and the (multi)set of n eigenvalues is called the* spectrum *of M.*

Here is a very useful fact about the trace:

Theorem. $\operatorname{tr}(M) = \sum_{i=1}^{n} \lambda_i$.

If you want to prove this, there's a set of problems that goes through a proof at the end of this handout.

Definition. The quantity

$$\rho(M) = \max_{1 \le i \le n} |\lambda_i|$$

is called the *spectral radius* of M.

1.3. Back to graphs

Unless otherwise noted, A is always the adjacency matrix of G.

The adjacency matrix of any graph is symmetric. (Why is that true?) So it has an eigenbasis! This will be very useful.

Problem 4. Find 3 linearly independent eigenvectors of the adjacency matrix for K_3. What are their eigenvalues?

Problem 5. Suppose that $v_1, \ldots, v_n$ are an eigenbasis for A with eigenvalues $\lambda_1, \ldots, \lambda_n$. Let $w_G(k)$ denote the number of closed walks of length k in G. Prove that

$$\boxed{\frac{w_G(k)}{\rho(A)^k} \longrightarrow 1}$$

as $k \to \infty$.

This answers our question very nicely! The number of closed walks of length k in any graph G is approximately $\rho(G)^k$, and $\rho(G)$ is something that you can get a computer to calculate fairly quickly.[1]

2. Miscellaneous facts

This is just the beginning of the connections between linear algebra properties of A and graph theory properties of G. Here are lots more!

To simplify things, we'll always assume that the eigenvalues are sorted so that

$$\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$$

[1] If you want to know *how* computers can calculate this, ask me at TAU.

and we'll write $\rho(G)$ sometimes instead of $\rho(A)$.

Problem 6. The *maximum degree* of G, denoted $\Delta(G)$, is the largest number of edges connected to any single vertex. Prove that $\rho(G) \leq \Delta(G)$. [Hint: If $Av = \lambda v$, how big can λ be?]

Problem 7. Prove that isomorphic graphs have the same spectrum.

Problem 8. A graph is *d-regular* if every vertex has degree d.
(a) Suppose that G is d-regular. Prove that $\lambda_n = d$. [Hint: Can you find an explicit eigenvector?]
(b) Suppose that $\lambda_n = \Delta(G)$ and G is connected. Prove that G is $\Delta(G)$-regular.

Problem 9. Prove that a d-regular graph is bipartite if and only if $\lambda_1 = -d$.

BONUS: MORE MISCELLANEA

Problem 10. A multiset is *symmetric* if λ appears the same number of times as $-\lambda$. Prove that G is bipartite if and only if the spectrum of G is symmetric.

Problem 11. If two graphs have the same spectrum, are they isomorphic?
(a) The star graph S_n consists of a single vertex connected to n other vertices (and no other edges). Find a basis of eigenvectors for the adjacency matrix of S_4.
(b) Find a basis of eigenvectors for the adjacency matrix of C_4 (the cycle with 4 vertices).
(c) Suppose that G consists of two connected components G_1 and G_2. Prove that the spectrum of G is the union of the spectrum of G_1 and the spectrum of G_2.
(d) Use the previous problems to find a pair of graphs that are not isomorphic but have the same spectrum.

Problem 12. The *complement* of G is the graph $\overline{G}$, which has the same vertices as G, but uv is an edge in $\overline{G}$ if and only if uv is *not* an edge in G. Suppose that G is d-regular. What is the spectrum of $\overline{G}$ in terms of the spectrum of G?

BONUS: TRACE AND EIGENVALUES

Problem 13.
(a) Prove that v is an eigenvector of M with eigenvalue λ if and only if v is an eigenvector of the matrix $\lambda I - M$ with eigenvalue 0.
(b) Prove that λ is an eigenvalue of M if and only if the rows of $\lambda I - M$ are linearly dependent.

Since $\det(M) = 0$ if and only if the rows of M are linearly dependent,[2] we can conclude that

Theorem. *λ is an eigenvalue of M if and only if* $\det(\lambda I - M) = 0$.

If you consider t as a variable and expand out the determinant of $tI - M$, you get a polynomial of degree t, called the *characteristic polynomial* of M, denoted $\chi_M(t)$. The roots of this polynomial are the eigenvalues of M.

Problem 14.
(a) Show that the coefficient of t^n in $\chi_M(t)$ is 1.
(b) Show that the coefficient of t^{n-1} in $\chi_M(t)$ is $\operatorname{tr}(M)$.
(c) Show that the constant term of $\chi_M(t)$ is $(-1)^n \det(M)$.

Problem 15. If M has an eigenbasis, then we know that $\chi_M(t)$ has roots $\lambda_1, \ldots, \lambda_n$. So in fact, $\chi_M(t) = (t - \lambda_1)(t - \lambda_2) \cdots (t - \lambda_n)$. Prove that $\det(M) = \prod_{i=1}^n \lambda_i$ and $\operatorname{tr}(M) = \sum_{i=1}^n \lambda_i$.

[2] If you haven't seen this before, we can talk at TAU!

Ghostly graph theory

Travis, Mathcamp 2024

3. diameter

Definition. The length of the smallest walk from x to y is called the *distance* from x to y and is denoted $d(x, y)$. The maximum distance between any pair of vertices is called the *diameter* of the graph, denoted $\operatorname{diam}(G)$.

Today, we'll see what the adjacency matrix can tell us about the diameter of a graph.

Problem 1. Prove that $d(x, y) \leq d(x, z) + d(z, y)$ for every triple of vertices x, y, z.

Definition. The set of all vertices whose distance from x is at most k is called the *ball of radius k* centered at x; the notation for this is $B_x(k)$.

Problem 2. Let Δ be the degree of G (that is, the maximum degree among all vertices in G). Find an upper bound on the size of $B_x(k)$ that depends only on k and Δ.

Problem 3. Prove that there is a constant c such that

$$\operatorname{diam}(G) \geq c \log_\Delta(n).$$

This makes some sense: If every vertex has small degree and the diameter is small, then there cannot be very many vertices.

Today, we'll use the adjacency matrix to find an *upper* bound on the diameter. Since $\rho(G) = d$ for any d-regular graph, let's define

$$\tilde{\rho} = \max_{1 \leq i \leq n-1} |\lambda_i|.$$

Theorem. *If G is a d-regular graph, then*

$$\operatorname{diam}(G) \leq \frac{\log(n)}{\log(d/\tilde{\rho})}.$$

Remember that the *dot product* of two vectors u and v is defined by

$$u \cdot v = u(1)v(1) + u(2)v(2) + \cdots + u(n)v(n),$$

where $u(i)$ is the ith coordinate of u.

To prove today's theorem, we'll need a stronger version of the spectral theorem:

Theorem (Spectral theorem, part 2). *If M is a symmetric matrix, then all of its eigenvalues are **real numbers** and it has an* orthonormal *eigenbasis. This means there are is a set of eigenvectors $v_1, v_2, \ldots, v_n$ that*

- *is a basis for $\mathbb{R}^n$, where*
- *$v_i \cdot v_i = 1$ (the length of each vector is 1), and*
- *$v_i \cdot v_j = 0$ if $i \neq j$ (the vectors are orthogonal).*

We won't prove the spectral theorem, but we'll use it a whole lot.

For the rest of today, assume G is a *connected* d-regular graph, and let $v_1, \ldots, v_n$ be an orthonormal eigenbasis of A, ordered so that the associated eigenvalues are increasing:

$$\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n = d.$$

Problem 4. Take $v_1, \dots, v_n$ as an orthonormal eigenbasis of A. What is v_n? (You can determine it (almost) exactly!)

Let e_x denote the vector that has entries

$$e_x(i) = \begin{cases} 1 & \text{if } i = x \\ 0 & \text{otherwise.} \end{cases}$$

Since $v_1, \dots, v_n$ is a basis of $\mathbb{R}^n$, there is a unique linear combination $e_x = \sum_{i=1}^n \alpha_i v_i$.

Problem 5.
(a) Find the values of α_n.
(b) What is $\sum_{i=1}^n \alpha_i^2$?

Problem 6. Prove: that $d(x, y) = \min\left\{k : e_x \cdot (A^k e_y) > 0\right\}$.

Problem 7. Prove today's main theorem. [Hint: How can you rewrite $e_x \cdot (A^k e_y)$? How can you guarantee this expression is positive?]

But there's a problem: What if $\tilde{\rho} = d$? Then the main theorem says that $\operatorname{diam}(G) \le \log(n)/\log(d/d) \dots$ which isn't a useful bound, oops.[3]

Problem 8. Suppose that G is a d-regular graph.
(a) Prove that the multiplicity of d as an eigenvalue of A is equal to the number of connected components of G.
(b) Suppose that G is bipartite. Prove that the multiplicity of $-d$ as an eigenvalue of A is equal to the number of connected components of G.

So if G is connected, then $\lambda_{n-1} < d$. But it still may be the case that $\tilde{\rho} = d$: maybe $\lambda_1 = -d$. We know that $\lambda_1 = -d$ if and only if G is bipartite. Fortunately, for bipartite graphs there is a special theorem. If G is a bipartite graph, define

$$\mu(G) = \max_{2 \le i \le n-1} |\lambda_i|.$$

Theorem. *If G is a connected d-regular bipartite graph, then*

$$\operatorname{diam}(G) \le \frac{\log(n/2)}{\log(d/\mu)} + 1.$$

Problem 9. Prove this theorem by modifying your previous proof.

If G is connected and bipartite, then $\mu(G) < d$, so this theorem gives a meaningful bound.

[3] Don't worry, there are staff supervising this class.

Ghostly graph theory

Travis, Mathcamp 2024

Today we have two more applications of linear algebra to graph theory.

4. mixing it up

Again, G is a d-regular graph. As usual, assume that $v_1, \ldots, v_n$ is an orthonormal eigenbasis for A, and $\lambda_1 \leq \cdots \leq \lambda_n$.

Problem 1. Given a subset $X \subseteq V(G)$, define $e_X \in \mathbb{R}^n$ by

$$e_X(i) = \begin{cases} 1 & \text{if } i \in X \\ 0 & \text{otherwise.} \end{cases}$$

Given $X, Y \subseteq V(G)$, suppose that $e_X = \sum_{i=1}^n \alpha_i v_i$ and $e_Y = \sum_{i=1}^n \beta_i v_i$.

(a) What is α_n?
(b) What is $\sum_{i=1}^n \alpha_i^2$?
(c) Express $e_X \cdot e_Y$ in terms of α_i and β_i. What does it mean combinatorially?

Definition. Given two vertex subsets $X, Y \in V(G)$, the number of edges from X to Y is

$$e(X,Y) = \left| \left\{ (x,y) \in X \times Y : xy \in E(G) \right\} \right|.$$

(Important! If $x, y \in X \cap Y$, then the edge xy is counted twice: once as (x,y) and once as (y,x).)

Problem 2. Suppose we had a set of vertices X and we built a graph as follows: for each vertex $x \in X$, choose d vertices independently at random in G, and connect x to each of them. (Multiple edges is okay.) Given another set $Y \subseteq V(G)$, what is the expected value of $e(X,Y)$?

The next theorem says that if the spectral radius is small, G mimics this random behavior:

Theorem (Expander mixing lemma). *For every $X, Y \subseteq V(G)$,*

$$\left| e(X,Y) - \frac{d|X||Y|}{n} \right| \leq \tilde{\rho}\sqrt{|X||Y|}.$$

Problem 3. Prove the theorem.

Problem 4. A set $X \subseteq V(G)$ is called *independent* if there is no edge between any pair of vertices in X. Prove that every independent set has size at most $\tilde{\rho}n/d$.

Problem 5. The *chromatic number* of a graph G is the minimum number of colors you need so that you can give each vertex a color in a way where no edge connects vertices of the same color. Prove that the chromatic number of G is at least $d/\tilde{\rho}$.

5. walking it off

In this section, do not assume that G is d-regular.

Problem 6. Prove that the number of walks of length k in G is at most $n\lambda_n^k$.

Problem 7. Let M be any symmetric matrix. Prove that

$$\rho(M) = \max \left\{ v \cdot (Mv) : v \cdot v = 1 \right\}.$$

(This is purely linear algebra.)

Now assume that G is a d-regular graph.

Problem 8. Let $X \subseteq V(G)$ and let $\gamma = |X|/n$. Suppose that $v \in \mathbb{R}^n$ satisfies the property that $v(i) = 0$ if $i \in X$ and $\|v\| = 1$, and decompose v according to the eigenbasis as $v = \sum_{i=1}^n \alpha_i v_i$.

(a) Prove that $\alpha_n^2 \leq 1 - \gamma$. [Hint: Express α_n as an inner product and use Cauchy-Schwarz.]

(b) Prove that $v \cdot (Av) \leq (1-\gamma)d + \gamma\tilde{\rho}$.

Problem 9. Prove that the number of walks of length k in a d-regular graph G that do not contain any vertex in X is at most

$$(1-\gamma)\big((1-\gamma)d + \gamma\tilde{\rho}\big)^k n.$$

This means that

Theorem. *The probability that a random walk of length k in a d-regular graph G completely avoids X is at most* $(1-\gamma)\big(1-\gamma(1-\frac{\tilde{\rho}}{d})\big)^k$.

Problem 10. Prove that theorem.

Ghostly graph theory

Travis, Mathcamp 2024

6. graph drawing

In this section, G is not necessarily d-regular.

Here is one of the coolest applications I know of spectral graph theory. The main question of this section is to figure out how to draw graphs in a way that makes sense.

Question. Given a very large graph G, how can we (or a computer) determine an informative way to draw it?

We'll think about this physically: Nature tends to optimize things that we don't know how to do, so imagine replacing each edge of the graph with a spring. The further the two endpoints of the edge are stretched apart, the more energy it takes to hold the endpoints in place. Specifically, the energy is proportional to the square of the distance between the endpoints. So if vertex i is at position $\xi(i)$, then the energy of a configuration is

$$\mathcal{E}(\xi) = \sum_{ij \in E(G)} \|\xi(i) - \xi(j)\|^2,$$

and a "stable configuration" is one that minimizes this quantity.

Trying to draw graphs in the plane will be a bit hard, so let's try drawing them on the plane first.

6.1. graph drawing on a line

We want to find a map $\xi\colon V(G) \to \mathbb{R}$ that minimizes $\mathcal{E}(\xi)$.

This is a good idea, but the configuration that minimizes the energy is the map $\xi\colon V(G) \to \mathbb{R}$ defined by $\xi(i) = 0$ for every i. :(

But that's a silly solution. So perhaps what we want to do is demand that $\|\xi\| = 1$.

Problem 1. Find a map $\xi\colon V(G) \to \mathbb{R}^2$ such that $\|\xi\| = 1$ and $\mathcal{E}(\xi) = 0$.

:(:(:(. Oh! But we forgot to say that the center of mass of the points should be at the origin, which we can write as $\xi \cdot \mathbf{1} = 0$. Then $\|\xi\| = 1$ really corresponds to fixing a scaling factor of the drawing. So here is our actual minimization problem:

$$\min_{\substack{\|\xi\|=1 \\ \xi\cdot\mathbf{1}=0}} \sum_{ij \in E(G)} |\xi(i) - \xi(j)|^2.$$

It turns out to optimize this, we'll need to use a *different matrix* associated to G.

Definition. The *Laplacian* of a graph G is an $n \times n$ matrix with entries

$$L_{i,j} = \begin{cases} \deg(i) & \text{if } i = j \\ -1 & \text{if } i \neq j \text{ and } ij \in E(G) \\ 0 & \text{otherwise.} \end{cases}$$

Since L is symmetric, it also has an orthonormal eigenbasis.

As usual, we let $v_1, \ldots, v_n$ be an orthonormal eigenbasis of A with eigenvalues $\lambda_1 \leq \cdots \leq \lambda_n$. We also let $u_1, \ldots, u_n$ be an orthonormal eigenbasis of L with eigenvalues $\mu_1 \leq \mu_2 \leq \cdots \leq \mu_n$.

Problem 2. Suppose in this problem that G is d-regular. How do the μ_i relate to the λ_i?

Problem 3. Prove that

$$x \cdot (Lx) = \sum_{ij \in E(G)} |x(i) - x(j)|^2$$

for every vector $x \in \mathbb{R}^n$.

Problem 4.

(a) Prove that every eigenvalue of L is nonnegative (even when G is not d-regular).

(b) What is u_1? What is μ_1? (Unlike the adjacency matrix, this is true even for non-regular graphs!)

Problem 5. What configuration ξ solves the minimization problem?

6.2. Back to the plane

Given a function $\xi\colon V(G) \to \mathbb{R}^2$, let R denote the $n \times 2$ matrix whose ith row is $\xi(i)$. If we define

$$\xi(i) = \big(x(i), y(i)\big),$$

then the first column of R is the vector x and the second column is the vector y.

Problem 6. Prove that

$$\operatorname{tr}(R^\top L R) = \sum_{ij \in E(G)} \big\|\xi(i) - \xi(j)\big\|^2$$

To get a good graph drawing in 2 dimensions, we want to minimize $\mathcal{E}(\xi)$ subject to the constraints

- $x \cdot 1 = 0$ and $y \cdot 1 = 0$, which implies that the centroid of the vertices is the origin; and
- $\|x\| = 1$ and $\|y\| = 1$, which prevents you from changing the energy just by scaling the picture.

But this isn't enough for a good drawing:

Problem 7. Find the configuration $\xi\colon V(G) \to \mathbb{R}^2$ that satisfies these constraints and minimizes $\mathcal{E}(\xi)$.

The problem is that the drawing you get from this configuration isn't really 2-dimensional—for this choice of ξ, all the vertices of G lie on a line! We can measure how 2-dimensional a configuration is by taking the dot product $x \cdot y$. If $|x \cdot y|$ is close to 1, then the points $\big(x(i), y(i)\big)$ are close to lying on the same line. (In statistics, this is called the *covariance*.) So to make it really a drawing in the *plane*, we will add the condition that $x \cdot y = 0$.

Problem 8 (Linear algebra lemma)**.** Since L is symmetric, the matrix $R^\top L R$ is also symmetric. So it has two real eigenvalues $\sigma_1 \le \sigma_2$. The point of this problem is to show that

$$\sigma_2 \ge \mu_3 \qquad \text{and} \qquad \sigma_1 \ge \mu_2.$$

(I recommend assuming this result and coming back to this problem if you have time.)

(a) Suppose that $x \cdot y = 0$ and $x \cdot \mathbf{1} = y \cdot \mathbf{1} = 0$. Show that

$$\{Ru : u \in \mathbb{R}^2\}$$

is a subspace of dimension 2 that is orthogonal to the vector $\mathbf{1}$.

(b) Show that $\|Ru\| = \|u\|$.

(c) Prove that there is a vector $u \in \mathbb{R}^2$ such that $Ru \in \operatorname{span}(u_3, u_4, \ldots, u_n)$.

(d) Why does this mean that $R^\top L R$ has an eigenvector with eigenvalue at least μ_3?

(e) Prove that the other eigenvector of $R^\top L R$ has eigenvalue at least μ_2.

Problem 9. Find a configuration $\xi\colon V(G) \to \mathbb{R}^2$ that minimizes $\mathcal{E}(\xi)$ among all those configurations that satisfy

- $x \cdot \mathbf{1} = y \cdot \mathbf{1} = 0$ and
- $\|x\| = \|y\| = 1$ and
- $x \cdot y = 0$.

Problem 10. What if you wanted to visualize a graph in 3 dimensions? What extra conditions would you need to add? What is the minimizing configuration?

BONUS: EFFICIENT RANDOM SEARCH ALGORITHMS

Using the last theorem from Day 3, we can make a more efficient random search algorithm!

Suppose that you have a set S of n elements and you want to find an element in a subset $X \subseteq S$ of size $|X| = cn$ (for $0 < c < 1$). But you don't know where the elements of X are; all you can do is test whether a given element of S is contained in X. So a reasonable strategy to find an element of X is to keep picking random elements of S.

Problem 11. To pick a random element of S, you need $\log(n)$ random bits. (Why?) If you pick a random element of S independently k times, how many random bits do you need? What is the probability that you find an element of X?

But random bits are hard to come by, and it would be nice if we could use fewer. It seems there's no hope to do better, but there is! Imagine that we can define a d-regular graph G whose vertex set is S. (Which there are algorithms for, though we won't talk about them.)

Problem 12. How many bits do you need to define a random walk of length k on the graph? What is the probability that the random walk visits a vertex in X?

Problem 13. Compare these two results.

BONUS: DIAMETER FOR NON-REGULAR GRAPHS

This has nothing to do with anything else on today's sheets.

Problem 14. In this problem, don't assume that G is regular.

(a) If $p(t)$ is a polynomial, say $p(t) = a_0 + a_1 t + \cdots + a_k t^k$, then we can evaluate p "at the matrix A" by writing $p(A) = a_0 + a_1 A + \cdots + a_k A^k$. Prove that there is no polynomial of degree $\leq \operatorname{diam}(G)$ such that $p(A) = \mathbf{0}$.[4]

(b) Suppose that A has exactly k distinct eigenvalues. Find a polynomial p of degree k such that $p(A) = \mathbf{0}$.

This proves:

> THEOREM. *If G is any connected graph with exactly k distinct eigenvalues, then* $\operatorname{diam}(G) \leq k - 1$.

[4] $\mathbf{0}$ is the $n \times n$ matrix where every entry is 0.

Ghostly graph theory

Travis, Mathcamp 2024

Today, G is a connected d-regular graph.

7. Eigenvalues and graph limits

7.1. Graph limits

Definition. A *rooted graph* is just a graph in which one vertex is marked as the root. If G is a rooted graph, $B_k(G)$ is the ball of radius k centered at the root of G.

Today, every graph is rooted and connected.
Also, graphs are allowed to have infinitely many vertices.

Using this, we can define a *distance* between two rooted graphs.

Definition. Suppose that G_1 and G_2 are two rooted graphs, and k is the maximal integer such that $B_k(G_1) = B_k(G_2)$. The *rooted distance* between G_1 and G_2 is $\delta(G_1, G_2) = 1/k$. (If there is no largest integer, then $\delta(G_1, G_2) = 0$.)

This is a metric:

Problem 1. Prove that

$$\delta(G_1, G_2) \le \max\big\{\delta(G_1, G_3), \delta(G_3, G_2)\big\}.$$

for every triple of rooted graphs G_1, G_2, G_3. (This is stronger than the usual triangle inequality.)

We can use this notion of rooted balls to define graph convergence.

Definition. We say that a sequence of rooted graphs $G_1, G_2, G_3, \ldots$ *converges to* a graph G if $\delta(G_n, G) \to 0$ as $n \to \infty$.

Problem 2. Suppose that $G_1, G_2, G_3, \ldots$ is a sequence of rooted graphs. Prove that there is a subsequence G_{n_i} such that for every $j \ge i \ge k$,

$$B_k(G_{n_i}) = B_k(G_{n_j}).$$

(In other words, the ball of radius k is the same in every G_{n_i} with $i \ge k$.)

Problem 3. Prove that every sequence of graphs has a convergent subsequence. (Remember that graphs are allowed to be countably infinite.)

7.2. Walks, yet again

Definition. Given a vertex $u \in V(G)$, we let $w_u(2k)$ denote the number of closed walks of length $2k$ that begin and end at u. If we need to specify the graph, we write $w_u^G(2k)$. We also define $\varpi_u(G) = \lim_{k\to\infty} w_u^G(2k)^{1/2k}$.

Problem 4. Prove that $\varpi_u(G) = \varpi_v(G)$ for every $u, v \in V(G)$, even if G is infinite. (So from now on, we'll just write $\varpi(G)$.)

Problem 5. Suppose that G is a finite graph. Show that for every k and every vertex u,

$$w_u^G(2k) \le \frac{n-1}{n}\tilde{\rho}^{2k} + \frac{d^{2k}}{n}.$$

Problem 6. Suppose that $G_1, G_2, G_3, \ldots$ is a sequence of graphs that converges to an infinite graph G. Prove that

$$\liminf_{n\to\infty} \tilde{\rho}(G_n) \geq \varpi(G).$$

(If you haven't seen "lim inf" before, this is just a way to get around the fact that the sequence $\tilde{\rho}(G_n)$ may not be convergent. You can just assume that the sequence $\tilde{\rho}(G_n)$ is convergent.)

7.3. A big theorem

Theorem. *If $G_1, G_2, G_3, \ldots$ is any sequence of connected d-regular graphs such that $|V(G_n)| \to \infty$, then*

$$\liminf_{n\to\infty} \tilde{\rho}(G_n) \geq 2\sqrt{d-1}.$$

Definition. The *infinite d-ary tree T_d* is the connected graph with countably many vertices and no cycles where every vertex has degree d.

For example, a small portion of the 3-ary tree looks like this:

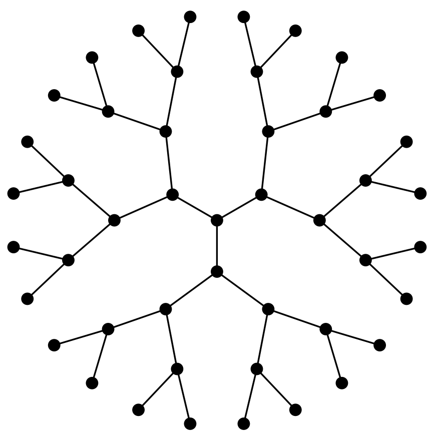

Problem 7. Prove that $w_u^G(2k) \geq w_v^{T_d}(2k)$ for any $u \in V(G)$ and $v \in V(T_d)$.

Here is a useful fact:

Proposition. $\varpi(T_d) = 2\sqrt{d-1}$.

That's all you need!

Problem 8. Prove the theorem.

Here is another way to state the theorem:

Theorem. *For every $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that every d-regular graph with at least N vertices has spectral radius $\tilde{\rho} \geq 2\sqrt{d-1} - \varepsilon$.*

In other words, any d-regular graph with many, many vertices must have a somewhat large spectral radius. Or, as another way to think about it: For every $\tilde{\rho} < 2\sqrt{d-1}$, there are only finitely many d-regular graphs whose spectral radius is at most $\tilde{\rho}$.

Bonus: walks in the d-ary tree

These problems prove that

$$\varpi(T_d) \geq 2\sqrt{d-1},$$

which is the inequality you need to prove the big theorem in the previous section.

Definition. The kth Catalan number 😺$_k$ the number of sequences of length $2k$ consisting of k $+$ symbols and k $-$ symbols such that: For every i, the first i symbols contain at least as many $+$'s as $-$'s. (For example, $++-+--$ is an allowable sequence, but $+--+-+$ is not.)

Problem 9. Prove that $w_u^{T_d}(2k) \geq (d-1)^k$ 😺$_k$.

The asymptotic notation $f(n) \sim g(n)$ means that $f(n)/g(n) \to 1$ as $n \to \infty$. (In other words, f and g grow at the same rate.)

Theorem. 😺$_k = \frac{1}{2k+1}\binom{2k}{k}$.

Theorem. $\binom{2k}{k} \sim \frac{4^k}{\sqrt{\pi k}}$.

If you want to know where these formulas come from, look up *Catalan numbers* and *Stirling's approximation* on Wikipedia.

Problem 10. Prove that $\varpi(T_d) \geq 2\sqrt{d-1}$.

Ghostly graph theory summary

Travis, Mathcamp 2024

Over the course of this class, you've proved a *lot* about graphs using linear algebra! Here are some of the highlights. For reference, $\rho(G) = \max_{1 \le i \le n} |\lambda_i|$ and

$$\tilde{\rho}(G) = \min_{1 \le i \le n-1} |\lambda_i|.$$

Theorem. *If G is a connected d-regular non-bipartite graph, then*

$$\frac{w_G(k)}{\rho(G)^k} \longrightarrow 1.$$

Theorem. *If G is a d-regular graph, then*

$$\operatorname{diam}(G) \le \frac{\log(n)}{\log(d/\tilde{\rho})}.$$

Theorem. *The chromatic number of a d-regular graph is at least $d/\tilde{\rho}$.*

Algorithm. *To plot a graph in a pleasing way, take the two eigenvectors u_2, u_3 of the Laplacian matrix and use them as coordinates for the vertices.*

Theorem. *If $G_1, G_2, G_3, \ldots$ is any sequence of connected d-regular graphs such that $|V(G_n)| \to \infty$, then*

$$\liminf_{n \to \infty} \tilde{\rho}(G_n) \ge 2\sqrt{d-1}.$$