Problem Set 4

Due: Wednesday, Oct. 12, email the pdf to toolkit2016homework@gmail.com

---

Homework policy: Exactly the same as last time.

Note: In this homework, the "Laplacian" always refers to the "random walk normalized Laplacian" operator, $L = I - K$, where K is the normalized adjacency matrix. There are also a couple of questions where you might find it easier to do the parts (a), (b) etc. in a different order; it's up to you.

---

1. Let $G = (V, E)$ be an undirected graph with Laplacian L. As discussed in class, there is always a complete basis $1 \equiv \phi_1, \phi_2, \dots, \phi_n$ of orthonormal eigenfunctions for L, with associated eigenvalues $\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$.

   (a) Let $\kappa_i = 1 - \lambda_i$ for $i \in [n]$. Show that κ_i is an eigenvalue of the normalized adjacency matrix K, with associated eigenvector ϕ_i. Deduce that $|\kappa_i| \le 1$ for all i.

   (b) Let G' be the "lazy" version of G, in which each vertex v gets $\deg(v)$ self-loops added to it. (Thus, the standard random walk on G' is like that on G, except that at each time step one lazily "does nothing" with probability $\frac{1}{2}$.) Show that the Laplacian L' of G' is $\frac{1}{2}L$. Defining $\lambda'_i = \frac{1}{2}\lambda_i$ for all $i \in [n]$, deduce that (λ'_i, ϕ_i) is an eigenvalue/eigenfunction pair for L', and that $0 \le \lambda'_i \le 1$.

   (c) Let G^2 denote the graph with vertex set V and with k edges between x and y if and only if there are k paths from x to y of length exactly 2 in G. Determine an explicit formula for the Laplacian of G^2 in terms of L, and also an explicit formula for its eigenvalues.

2. (a) Let G be the complete graph on n vertices, with self-loops. (This is an n-regular graph.) Determine the eigenvalues $\lambda_1, \dots, \lambda_n$ of G's Laplacian.

   (b) Let G be the "hypercube graph" with vertex set $V = \{-1, +1\}^d$ and an edge connecting two strings if they differ on exactly one coordinate. For each set $S \subseteq [d]$, define $\phi_S : V \to \mathbb{R}$ by $\phi_S(u) = \prod_{i \in S} u_i$. Show that each ϕ_S is an eigenfunction for G's Laplacian L, and determine all eigenvalues of L.

3. Let $G = (V, E)$ be an undirected graph. The (NP-hard) *Max-Cut* problem is to partition V into two parts S, $V \setminus S$ so as to maximize the fraction of "cut" edges $e \in E$, those with one endpoint in S and one in $V \setminus S$. Let $\mathrm{Opt}(G) \in [0, 1]$ denote the largest fraction of edges in G that can be cut.

   (a) Show that $\mathrm{Opt}(G) \le \frac{1}{2}\lambda_n(L)$, where $\lambda_n(L)$ denotes the largest eigenvalue of G's Laplacian L. (Hint: consider $f : V \to \{\pm 1\}$.)

   (b) More generally, show that $\mathrm{Opt}(G) \le \frac{1}{2}\lambda_n(L')$ for every matrix L' of the form $L + \mathrm{diag}(h)$, where $h : V \to \mathbb{R}$ satisfies $\mathbf{E}_{\boldsymbol{u} \sim \pi}[h(\boldsymbol{u})] = 0$ (and where $\mathrm{diag}(h)$ denotes the diagonal matrix with "vector" h on its diagonal).
   (Remark: The smallest possible value of $\frac{1}{2}\lambda_n(L')$, over all valid choices of h, is equal to the "Goemans–Williamson SDP relaxation" value of the Max-Cut. We will find out more about this in a later lecture.)

4. This problem uses the notation from Problem 1(a).

   (a) Suppose that, for $f : V \to \mathbb{R}$, we introduce the notation $\widehat{f}(i) = \langle f, \phi_i \rangle$ (where, as usual, $\langle g, h \rangle = \mathbf{E}_{\boldsymbol{u}\sim\pi}[g(\boldsymbol{u})h(\boldsymbol{u})]$ and π is the invariant distribution for G). Show that

   $$\mathop{\mathbf{E}}_{\boldsymbol{u}\sim\boldsymbol{v}}[f(\boldsymbol{u})g(\boldsymbol{v})] = \sum_{i=1}^{n} \kappa_i \widehat{f}(i)\widehat{g}(i).$$

   (For the next part, we recommend you recall, from class, how to express $\mathbf{E}[f]$ and $\mathbf{Var}[f]$ in terms of the $\widehat{f}(i)$'s.)

   (b) Let $S, T \subseteq V$, and recall the notation $\mathrm{vol}(S) = \mathbf{Pr}_{\boldsymbol{u}\sim\pi}[\boldsymbol{u} \in S]$. Prove that

   $$\left|\mathop{\mathbf{Pr}}_{\boldsymbol{u}\sim\boldsymbol{v}}[\boldsymbol{u} \in S, \boldsymbol{v} \in T] - \mathrm{vol}(S)\mathrm{vol}(T)\right| \leq \kappa\sqrt{\mathrm{vol}(S)(1-\mathrm{vol}(S))\mathrm{vol}(T)(1-\mathrm{vol}(T))},$$

   where $\kappa := \max_{i\neq 1}|\kappa_i| = \max\{|\kappa_2|, |\kappa_n|\}$. Naturally, you will need to use the Cauchy–Schwarz inequality.

   (Remark: if we put the weaker bound $\sqrt{\mathrm{vol}(S)\mathrm{vol}(T)}$ on the right-hand side above, the resulting inequality is sometimes called the "Expander Mixing Lemma".)

5. Recall from Problem 1(b) the technique of making a graph G "lazy". One reason this is often done is to get rid of "even-odd time issues": a lazy graph is never bipartite, so a "lazy random walk" on a connected graph G always converges to the invariant distribution π (albeit in about twice the time). Also, as you saw in that problem, the Laplacian eigenvalues of a lazy graph are always between 0 and 1, and hence the normalized adjacency eigenvalues are always nonnegative; this too is quite pleasant. The present problem describes an elegant alternative to "making G lazy"; it also fixes the issues of bipartite graphs / negative eigenvalues of K / even-odd time.

   Let $0 < \delta \leq 1$. Imagine we do a random walk on graph G starting at time 0. At each time instant δ, 2δ, 3δ, etc., we do the following: with probability δ, take a random step in the graph; with probability $1-\delta$, stay put.

   Now for any fixed $t \in \mathbb{R}^{\geq 0}$, we will consider the position of the walk at instant t. (If t/δ is an integer, we include the possible step made at instant t.) Note that the $\delta = 1$ case is just the "standard random walk", and the $\delta = 1/2$ case is the "lazy random walk" (with time "sped up" by a factor of 2).

   The other elegant value of δ to consider is "infinitesimally small". This case is called the "continuous-time random walk on G". Note that for a fixed δ, the number of graph-steps taken by time t is distributed as $\boldsymbol{T} \sim \mathrm{Binomial}(t/\delta, \delta)$, a random variable with mean t. You may know that as $\delta \to 0$, this random variable converges to having the "Poisson(t)" distribution with mean t. This is defined by

   $$\mathbf{Pr}[\boldsymbol{T} = s] = e^{-t}\frac{t^s}{s!} \qquad \text{for each } s \in \mathbb{N}. \tag{1}$$

   You can now forget about "δ" if you want, and just use the following definition: "Do a continuous-time random walk on G for time t" means "Take $\boldsymbol{T}$ random walk steps in G, where $\boldsymbol{T}$ is a random variable having the Poisson(t) distribution".[1]

[1] An alternate definition: Notice that the waiting time between successive random walk steps is $\delta\boldsymbol{S}$, where

(a) Define the "continuous-time transition operator (AKA heat operator)" H_t on functions as follows: $H_t f(u) = \mathbf{E}[f(\boldsymbol{v})]$, where $\boldsymbol{v}$ is the (random) vertex gotten by doing a continuous-time random walk in G for time t, starting from vertex u. (You should think of H_t as being roughly analogous to K^t, at least when t is an integer.) Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of G's Laplacian L and $\phi_1, \dots, \phi_n$ the associated eigenfuctions, as usual. Show that $\phi_1, \dots, \phi_n$ are also eigenfunctions of H_t, with associated eigenvalues $1 = \exp(-t\lambda_1) \geq \exp(-t\lambda_2) \geq \cdots \geq \exp(-t\lambda_n) \geq 0$.

Remark: For your own edification, I strongly encourage you to plot, say $\exp(-5\lambda)$ and $(1-\lambda)^5$ (the eigenvalues of H_5 and K^5) versus λ, for $\lambda \in [0, 2]$.

(b) Show that we can write H_t as the "matrix exponential" $\exp(-tL)$. Here the definition of $\exp(B)$ when B is a matrix is the entry-wise limit $I + B + \frac{1}{2!}B^2 + \frac{1}{3!}B^3 + \cdots$.

6. Let $G = (V_G, E_G)$ be an n-vertex, D-regular graph and let $H = (V_H, E_H)$ be a D-vertex, d-regular graph. Assume that each vertex in G has its edges explicitly labeled as 1, 2, ..., D (so that we may speak of the "ith neighbor of vertex $x \in V_G$" for $i \in [D]$). Also, assume the vertices V_H are identified with $[D]$. The *balanced replacement product* of G and H, denoted $G \textcircled{r} H$ is the $2d$-regular graph with vertex set $V_G \times V_H$ in which each vertex of G is replaced by a copy of H, and in addition d parallel edges are put between vertices (x, i) and (x', i') when x' is the ith neighbor of x in G and x is the i'th neighbor of x' in G. Here is a picture when G is the 3-dimensional hypercube graph and H is the triangle:[2]

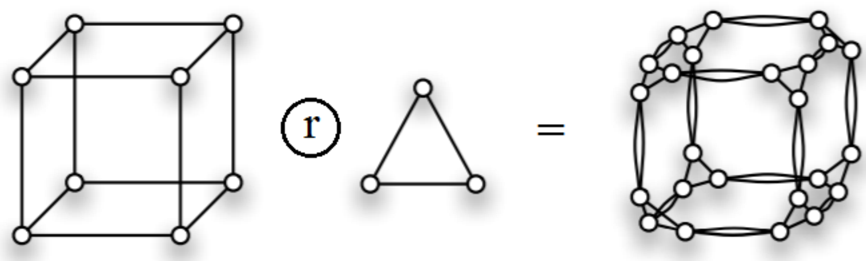

Let λ_G (respectively, λ_H) be the second-smallest eigenvalue of G's (respectively, H's) Laplacian. Show that the second-smallest eigenvalue, λ_2, of $G \textcircled{r} H$'s Laplacian satisfies $\lambda_2 \geq \frac{\lambda_G \lambda_H}{c}$ for some universal constant c.

(Hint: We recommend you begin by recalling why it suffices to compare $\mathbf{Var}[f]$ and $\mathcal{E}[f]$ for $f : V_G \times V_H \to \mathbb{R}$. Then we suggest the following lines: "For each $x \in V_G$, write $f_x : V_H \to \mathbb{R}$ for the function defined by $f_x(i) = f(x, i)$. Also, define $\overline{f} : V_G \to \mathbb{R}$ by $\overline{f}(x) = \mathbf{E}_{\boldsymbol{i} \sim V_H}[f(x, \boldsymbol{i})]$. By the the Law of Total Variance,

$$\mathop{\mathbf{Var}}_{(\boldsymbol{x},\boldsymbol{i}) \sim V_G \times V_H}[f(\boldsymbol{x}, \boldsymbol{i})] = \mathop{\mathbf{E}}_{\boldsymbol{x} \sim V_G}\left[\mathop{\mathbf{Var}}_{\boldsymbol{i} \sim V_H}[f_{\boldsymbol{x}}(\boldsymbol{i})]\right] + \mathop{\mathbf{Var}}_{\boldsymbol{x} \sim V_G}\left[\overline{f}(\boldsymbol{x})\right].$$

Now..."

One more hint: the "approximate triangle inequality" $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$.)

$\boldsymbol{S} \sim$ Geometric(δ), a random variable with mean 1. As $\delta \to 0$, this random variable converges in distribution to the exponential "Exp(1)" random variable studied in Problem 2 of Homework 3. An equivalent definition of the "continuous-time random walk on G" is "wait for an Exp(1) amount of time, then do a random step in G; then wait for another (independent) Exp(1) amount of time, then do a random step in G; then wait for another (independent) Exp(1) amount of time", etc.

[2]Thanks to Petteri Kaski for drawing most of this figure.