Spectral Graph Theory

R. Kyng & M. Probst *Problem Set 4 — Monday, March 11*

These exercises will not count toward your grade, but you are encouraged to solve them all. This exercise sheet contains exercises relating to lectures in Week 4.

To get feedback, you must hand in your solutions by 23:59 on March 21. Both hand-written and LaTeX solutions are acceptable, but we will only attempt to read legible text.

Exercise 1

The bound obtained in Cheeger's inequality is indeed tight. Prove that:

1. Let G be the graph consisting of two vertices connected by a single edge of unit weight. Prove that $\phi(G) = \lambda_2(\boldsymbol{N})/2$ and therefore that the lower bound of Cheeger's inequality is tight.
2. To show that the line graph proves that the upper bound of Cheeger's Inequality is asymptotically tight (i.e. up to constant factors).

Exercise 2

Sparse Expanders: In random graph theory, the graph over n vertices where each edge between two endpoints is present independently with probability p is denoted $G(n, p)$.

Show that for $p = \Omega(\log n/n)$, that $G(n, p)$ is a $\Omega(1)$-expander with high probability (it is up to you to fix large constants). Take the following steps:

1. Prove that with high probability, $\boldsymbol{d}(u) = \Theta(pn)$ for all vertices $u \in V(G(n, p))$.
2. For each set S of $k \leq n/2$ vertices, argue that
$$\mathbb{P}[|E(S, V \setminus S)| = \Theta(kpn)] > 1 - n^{-c \cdot k}$$
for any large constant $c > 0$.
3. Observe that there are at most $\binom{n}{k}$ sets of vertices S of size k. Conclude that $G(n, p)$ is with high probability a $\Omega(1)$-expander.

Exercise 3

Let $G = (V, E)$ be a connected, undirected graph. In this problem, you will show that there is an algorithm that computes a ϕ-expander decomposition $X_1, X_2, \ldots, X_k$ for G of quality $q = O(\phi^{-1/2} \cdot \log n)$ in time $O(m \log^c n)$ for some constant c. We let $\boldsymbol{N}$ denote the normalized Laplacian in this exercise, defined by $\boldsymbol{N} = \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2}$. Assume that you are given an algorithm $\textsc{CertifyOrCut}(G, \phi)$ that given a graph G and a parameter ϕ either:

- *Certifies* that G is a ϕ-expander, or
- Presents a *cut* S such that $\phi(S) = O(\sqrt{\phi})$.

The algorithm $\textsc{CertifyOrCut}(G, \phi)$ runs in time $O(m \log^{c'} n)$ for $c' > 0$.

1. Show that there is an algorithm that uses $\textsc{CertifyOrCut}(G, \phi)$ and computes a ϕ-expander decomposition of quality $O(\phi^{-1/2} \cdot \log n)$ in time $O(mn \cdot \log^{c'} n)$.
2. Show that in $O(mn \cdot \log^{c'} n)$ time, you can implement a procedure $\textsc{CertifyOrLargeCut}(G, \phi)$ that outputs a set S (possibly empty) with $\phi(S) = O(\sqrt{\phi})$ such that either
   - $G[V \setminus S]$ is a ϕ-expander and $vol_G(V \setminus S) \geq \frac{1}{3}m$, or
   - $\min\{vol_G(S), vol_G(V \setminus S)\} \geq \frac{1}{3}m$.

   Comment: You may use the following claim.

   Claim: Given a set $S \subseteq V$ of conductance $\phi_G(S) \leq \phi$ and a set $S' \subseteq V \setminus S$ in $G[V \setminus S]$ with conductance $\phi_{G[V \setminus S]}(S') \leq \phi$ such that $vol_G(S \cup S') \leq vol_G(V)/2$, then $\phi_G(S \cup S') \leq \phi$.
3. **(BONUS)** Using a local version of random walks and heavy randomization, it turns out that $\textsc{CertifyOrLargeCut}(G, \phi)$ can be implemented in $O(m \log^{c''} n)$ time[1]. Show that this implies an $O(m \log^{c''+1} n)$ time algorithm to compute a ϕ-expander decomposition.

Exercise 4

For a matrix $\boldsymbol{Z}$ to be the pseudoinverse of a symmetric matrix $\boldsymbol{M}$, you need to show that

1. $\boldsymbol{Z}^\top = \boldsymbol{Z}$.
2. $\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{0}$ for $\boldsymbol{v} \in \ker(\boldsymbol{M})$.
3. $\boldsymbol{M}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{v}$ for $\boldsymbol{v} \in \ker(\boldsymbol{M})^\perp$.

Prove that if $\boldsymbol{Z}$ and $\boldsymbol{Y}$ are both the pseudo-inverse of symmetric matrix $\boldsymbol{M}$, then $\boldsymbol{Z} = \boldsymbol{Y}$, i.e. the pseudo-inverse is unique.

Exercise 5

Let $\boldsymbol{M} = \boldsymbol{X}\boldsymbol{Y}\boldsymbol{X}^\top$ for some $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{R}^{n \times n}$, where $\boldsymbol{X}$ is invertible and $\boldsymbol{M}$ is symmetric. Furthermore, consider the spectral decomposition of $\boldsymbol{M} = \sum_{i=1}^n \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$. Then, we define $\boldsymbol{\Pi}_{\boldsymbol{M}} = \sum_{i, \lambda_i \neq 0} \boldsymbol{v}_i \boldsymbol{v}_i^\top$. $\boldsymbol{\Pi}_{\boldsymbol{M}}$ is the orthogonal projection onto the image of $\boldsymbol{M}$: It has the property that for $\boldsymbol{v} \in \mathrm{im}(\boldsymbol{M})$, $\boldsymbol{\Pi}_{\boldsymbol{M}}\boldsymbol{v} = \boldsymbol{v}$ and for $\boldsymbol{v} \in \ker(\boldsymbol{M})$, $\boldsymbol{\Pi}_{\boldsymbol{M}}\boldsymbol{v} = \boldsymbol{0}$.

Prove that

$$\boldsymbol{Z} = \boldsymbol{\Pi}_{\boldsymbol{M}} (\boldsymbol{X}^\top)^{-1} \boldsymbol{Y}^+ \boldsymbol{X}^{-1} \boldsymbol{\Pi}_{\boldsymbol{M}}$$

is the pseudoinverse of $\boldsymbol{M}$.

[1] Note that this is slightly idealized for ease of presentation.