Toolkit for TCS		Feb-June 2021
	HW : Spectral Methods and Polynomial Methods	
Out: June 2, 2021		**Due:** June 16, 2021

(1) Please take time to write clear and concise solutions. You are STRONGLY encouraged to submit LaTeXed solutions.
(2) Collaboration is OK, but please write your answers yourself, and include in your answers the names of EVERYONE you collaborated with and ALL references other than class notes you consulted.
(3) Answer any 4 of the 5 questions. Each question has the same weightage of 10 points each.

1. **[Trevisan's robust characterization of bipartiteness]**

   Let $G = (V, E)$ be an undirected unweighted connected graph. Let W be the random-walk matrix and $L = I - W$ be the Laplacian. Recall that $W = D^{-1}A$ where $D = \mathrm{Diag}(deg)$ is the diagonal matrix of degrees and A is the adjacency matrix of G. Let π be the stationary distribution and $\langle \cdot, \cdot \rangle_\pi$ the corresponding π-inner product. Recall that the quadratic form corresponding to the Laplacian satisfies the following.

$$\langle f, Lf \rangle_\pi = \sum_{\{i,j\}} \pi(i) \cdot W(i,j) \cdot (f(i) - f(j))^2,$$

   where π the stationary distribution of this random walk matrix The largest eigenvalue of the normalized Laplacian, denoted by λ_n, satisfies

$$\lambda_n = \max_{f \neq 0} \frac{\langle f, Lf \rangle_\pi}{\langle f, f \rangle_\pi}.$$

   (a) (1 point) **[bipartite $\Leftrightarrow$ $\lambda_n = 2$]**
   Prove that $\lambda_n \leq 2$. Furthermore, prove that equality holds iff the graph G is bipartite.

   (b) (2 points) **[almost bipartite $\Rightarrow$ λ_n almost 2]**
   Suppose the MAXCUT in G has normalized cost at least $1 - \varepsilon$. That is, there exists a cut $(S, V \setminus S)$ such $|E(S, V \setminus S)| \geq (1 - \varepsilon)|E|$ where $E(S, V \setminus S) = \{\{u, v\} \in E : u \in S, v \notin S\}$. Prove that there is a non-zero vector $f : V \to R$ such that

$$\langle f, Lf \rangle_\pi \geq \frac{1 - \varepsilon}{2},$$
$$\langle f, f \rangle_\pi = \pi(S) \cdot (1 - \pi(S)) \leq \frac{1}{4}.$$

   Hence, conclude that $\lambda_n \geq 2(1 - \varepsilon)$.

   (c) (7 points) **[λ_n almost 2 $\Rightarrow$ almost bipartite]**
   In this part, we will prove the following theorem.

   Theorem. *Let $\lambda_n \geq 2(1 - \varepsilon)$ or equivalently there exists a non-zero vector $f : V \to \mathbb{R}$ such that $\langle f, (I + W)f \rangle_\pi \leq 2\varepsilon \cdot \langle f, f \rangle_\pi$. Then there exists non-zero vector $y \in \{-1, 0, 1\}^V$ such that*

$$\frac{\sum_{\{i,j\} \in E} |y_i + y_j|}{\sum_{i \in V} d_i |y_i|} \leq \sqrt{8\varepsilon}.$$

   To this end, we define the following randomized process that constructs a random non-zero vector $Y \in \{-1, 0, 1\}^V$ given a non-zero vector $f : V \to \mathbb{R}$ satisfying $\langle f, (I + W)f \rangle_\pi \leq 2\varepsilon \cdot \langle f, f \rangle_\pi$. Since this latter condition is scale-invariant, we may assume wlog. that $\max_i |f(i)| = 1$ and let $i_* \in V$ such that $|f(i_*)| = 1$.

   - Pick a value t uniformly in $[0, 1]$.
   - Define $Y \in \{-1, 0, 1\}^V$ as follows:

$$Y_i = \begin{cases} -1 & \text{if } f(i) < -\sqrt{t}, \\ 1 & \text{if } f(i) > \sqrt{t}, \\ 0 & \text{otherwise, i.e., } |f(i)| \leq \sqrt{t}. \end{cases}$$

i. (1 point) Prove that $\mathbf{P}[\exists i \in V, Y_i \neq 0] = 1$.
ii. (2 points) Prove that for any $i, j \in V$, $\mathbf{E}[|Y_i|] = f(i)^2$ and $\mathbf{E}\left[|Y_i + Y_j|\right] \leq |f(i) + f(j)| \cdot (|f(i)| + |f(j)|)$.
iii. (3 points) Prove that $\mathbf{E}\left[\sum_{\{i,j\}\in E} |Y_i + Y_j|\right] \leq \sqrt{8\varepsilon} \cdot \mathbf{E}\left[\sum_i d_i |Y_i|\right]$.
Hint: Cauchy-Schwarz Inequality.
iv. (1 point) Hence, conclude that there exists a non-zero vector $y \in \{-1, 0, 1\}^V$ such that $\sum_{\{i,j\}\in E} |y_i + y_j| \leq \sqrt{8\varepsilon}. \cdot \sum_{i \in V} d_i |y_i|$.

Discussion. *It is known that G is connected if $\lambda_2 \neq 0$. Or equivalently, $\phi(G) \neq 0$ iff $\lambda_2 \neq 0$. Cheeger's inequalities give a "quantitative strengthening" of this statement by showing that*

$$\sqrt{2\lambda_2} \geq \phi(G) \geq \lambda_2/2.$$

The problem is similar in spirit but works with λ_n and "bipartiteness" instead of λ_2 and "connectedness".

Define the bipartiteness ratio number *of a graph G to be*

$$\beta(G) := \min_{y \in \{-1,0,1\}^V} \frac{\sum_{\{u,v\}\in E} |y_u + y_v|}{2d \sum_{u \in V} |y_u|},$$

which is equivalent to

$$\beta(G) = \min_{S \subseteq V, (L,R) \text{ partition of } S} \frac{2\partial(L, L) + 2\partial(R, R) + \partial(S, V \setminus S)}{d|S|},$$

Observe that $\beta(G) = 0$ iff G is bipartite. Problem 1a shows that $\beta(G) = 0$ iff $\lambda_n = 2$. Problems 1b–1c are a quantitative strengthening of this claim as they demonstrate that

$$\sqrt{2(2 - \lambda_n)} \geq \beta(G) \geq \frac{1}{2} \cdot (2 - \lambda_n).$$

This result is due to Luca Trevisan.

2. **[Chernoff bound for Expander Random-Walks]**

In lecture, we proved the following hitting set lemma for random-walks on spectral expanders.

Lemma. *Let W be a reversible random walk on a set V of n vertices and $1 = \omega_1 \geq \omega_2 \geq \cdot \geq \omega_n \geq -1$ be its eigenvalues and π the stationary distribution. Let $\omega = \max\{\omega_2, |\omega_n|\}$. Let $B \subseteq V$ such that $\mu := \pi(B)$. Let $X_1, \ldots, X_t$ be a random-walk of length t according to W where the first vertex X_1 is chosen according to the stationary distribution. Then*

$$\mathbf{P}\left[\bigwedge_{i\in[t]} (X_i \in B)\right] \leq \mu \cdot (\mu + \omega(1 - \mu))^{t-1}. \quad (1)$$

In this problem, we will extend this obtain the following Chernoff-like bound on expander random-walks.

Theorem. *Let W be a reversible random walk on a set V of n vertices and $1 = \omega_1 \geq \omega_2 \geq \cdot \geq \omega_n \geq -1$ be its eigenvalues and π the stationary distribution. Let $\omega = \max\{\omega_2, |\omega_n|\}$. Let $B \subseteq V$ such that $\mu := \pi(B)$. Let $X_1, \ldots, X_t$ be a random-walk of length t according to W where the first vertex X_1 is chosen according to the stationary distribution. Then for any $\delta \in (0, 1)$, we have*

$$\mathbf{P}\left[\#\{i \colon X_i \in B\} \geq (\mu + \omega(1 - \lambda) + \delta)t\right] \leq \exp\left(-\Omega(\delta^2 t)\right).$$

Let $S \subseteq [t]$ be a random subset chosen as follows: for each $i \in [t]$, independently add i to S with probability q. This satisfies that for any fixed set $s \subseteq [t]$, we have $\mathbf{P}[S = s] = q^k \cdot (1 - q)^{t-k}$ where $k = |s|$.

(a) (2 points) Use (1) to conclude that for each integer $0 \le k \le t$

$$\mathbf{P}_{X_1,\cdots,X_t,S}\left[\bigwedge_{i\in S}(X_i \in B) \mid |S| = k\right] \le \mu\cdot(\mu+\omega(1-\mu))^{k-1} \le (\mu+\omega(1-\mu))^k.$$

Note that the above probability is over the random choice of the walk $X_1,\ldots,X_t$ as well as the set S conditioned on the fact that $|S| = k$.

(b) (4 points) Show that

$$\mathbf{P}_{X_1,\cdots,X_t,S}\left[\bigwedge_{i\in S}(X_i \in B)\right] \le (q\cdot(\mu+\omega(1-\mu))+1-q)^t.$$

(c) (3 points) Let X_B be the random subset of $[t]$ defined as follows:

$$X_B = \{i \in [t] : X_i \in B\}.$$

Show that

$$\mathbf{P}\left[|X_B| \ge (\mu+\varepsilon)t\right] \le \left(\frac{q\cdot(\mu+\omega(1-\mu))+1-q}{(1-q)^{1-\mu-\varepsilon}}\right)^t.$$

(d) (1 point) Let $\varepsilon > \lambda(1-\mu)$. Use calculus to show that the right hand side of the above expression is minimized when

$$q = \frac{\varepsilon-\omega(1-\mu)}{(1-\mu-\omega(1-\mu))\cdot(\mu+\varepsilon)},$$

to obtain

$$\begin{aligned}\mathbf{P}\left[|X_B| \ge (\mu+\varepsilon)t\right] &\le \left[\left(\frac{mu+\omega(1-\mu)}{\mu+\varepsilon}\right)^{\mu+\varepsilon}\cdot\left(\frac{1-\mu-\omega(1-\mu)}{1-\mu-\varepsilon}\right)^{1-\mu-\varepsilon}\right]^t\\ &= \exp\left(-D_{KL}(\mu+\varepsilon\|\mu+\omega(1-\mu))\cdot t\right).\end{aligned}$$

Discussion. *This proof of the Chernoff bound on expander random-walks is due to Impagliazzo and Kabanets. A proof can of the standard Chernoff bound can also be obtained along similar lines. For expander random-walks, a stronger Chernoff bound is known due to Gillman.*

3. **[SDP formulation of the Lovász Theta function]**

In class, we discussed the followin SDP formulation of the Lovász Theta function. follows

$$\theta_L(G) = \min_{\text{symmetric } M\in\mathbb{R}^{V\times V},\lambda} \lambda$$

subject to

$$\begin{cases} M_{u,v} = 1 & \text{if } u = v \text{ or } (u,v) \notin E \\ M \preccurlyeq \lambda I \end{cases}$$

In this problem, we will show that $\theta_L(G) = \vartheta(G)$.

(a) (5 points) Given an orthonormal representation $(u_1,\ldots,u_n)$ and a handle c with value $\vartheta(G,\overline{u},c)$, define the vectors $v_i := c - u_i/\langle c,u_i\rangle$. Consider the matrix $N = (\langle v_i,v_j\rangle)_{i,j}$ and $D = diag(\vartheta(G,\overline{u},c) - 1/\langle c,u_i\rangle^2)$. Use the above to show that $N + D \succcurlyeq 0$. Conclude that $\theta_L(G) \le \vartheta(G)$.

(b) (5 points) Let (M,λ) be a feasible solution to the SDP formulation of $\theta_L(G)$. Since $\lambda I - M \succcurlyeq 0$, we have that there exist vectors $v_1,\ldots,v_n$ such that $\lambda I - M = (\langle v_i,v_j\rangle)_{i,j}$. Let c be any unit vector orthogonal to all the v_i's. Define $u_i := (c+v_i)/\sqrt{\lambda}$. Show that $(u_1,\ldots,u_n)$ is an orthonormal representation of G. Using this prove that $\vartheta(G) \le \lambda$. Hence, conclude that $\vartheta(G) \le \theta_L(G)$.

4. **[Kakeya Sets]**

Let $\mathbb{F}_q$ be a finite field of size q. A *Kakeya set* in $\mathbb{F}_q^m$ is a set $K \subseteq \mathbb{F}_q^n$ such that K contains a line in every direction. More precisely, K is a Kakeya set if for every $y \in \mathbb{F}_q^m$ there exists a $z \in \mathbb{F}_q^m$ such that the line

$$L_{z,y} = \{z + t \cdot y \colon t \in \mathbb{F}_q\}$$

is contained in K.

A trivial upper bound on th size of K is q^m and this can be improved to $q^m/2^{m-1}$. In this problem, we will use the polynomial method to show a lower bound of $q^m/m!$. More precisely, we will show that

$$|K| \geq \binom{q+m-1}{m}.$$

Suppose, for contradiction that this is not the case.

(a) (4 points) Show that there exists a m-variate non-zero polynomial g of total degree $d \leq q-1$ such that $g(x) = 0$ for all $x \in K$.

Let g_d be the homogenous part of degree d of g so that g_d is non-zero and homogenous.

For any $y \in \mathbb{F}_q^m$, we know that there exists a $z \in \mathbb{F}_q^m$ such that the line $L_{z,y}$ is contained in K. Consider the following univariate polynomial

$$P_{y,z}(t) := g(z + t \cdot y).$$

(b) (2 points) Argue that $P_{y,z}$ is identically zero and hence the coefficient of t^d in $P_{y,z}(t)$ is zero.

(c) (2 points) Show that the coefficient of t^d in $P_{y,z}(t)$ is exactly $g_d(y)$.

(d) (2 points) Conclude that g_d is identically zero, a contradiction.

Discussion. *This proof is due to Zeev Dvir.*

5. **[3-AP-free sets in $\mathbb{F}_3^n$ via the polynomial method]**

Let $A \subseteq \mathbb{F}_3^n$. We say that A is 3-AP-free if there does not exist $x \neq y \in \mathbb{F}_3^n$ such that $x, (x+y)/2, y \in A$ (i.e., A does not contain any non-trivial arithmetic progression of length 3). In this problem, we will use the polynomial method to show that any 3-AP-free set is of size at most c^n for some fixed $c \in (2,3)$.

(a) (2 points) Let $0 \leq d \leq 2n$. Let $V_d(n)$ denote the set of all functions from $\mathbb{F}_3^n$ to $\mathbb{F}_3$ expressible as degree d polynomials. In other words, if $f \in V_d$, then f can be expressed as a polynomial of the form

$$f(x_1, \ldots, x_n) = \sum_{\mathrm{a}=(a_1,\ldots,a_n)\in\{0,1,2\}^n : \sum a_i \leq d} c_{\mathrm{a}} \prod_{i=1}^{n} x_i^{a_i}.$$

Let $m_d(n) = \dim(V_d(n))$. Prove the following facts about m_d.

i. $m_{2n}(n) = 3^n$.

ii. For all $0 \leq d \leq 2n$, $m_{2n-d}(n) = 3^n - m_d(n)$.

iii. There exists $c \in (2,3)$ such that $m_{2n/3}(n) \leq c^n$.

(b) (3 points) Let $A \subseteq \mathbb{F}_3^n$ be 3-AP-free.

i. Show that if $m_d > 3^n - |A|$, then there exists a non-zero $f \in V_d$ such that $f(\mathbf{x}) = 0$ for all $\mathbf{x} \notin A$.

ii. Strengthen the above to show that if $m_d > 3^n - |A|$, then there exists an $f \in V_d$ such that $f(\mathbf{x}) = 0$ for all $\mathbf{x} \notin A$ and f is non-zero on at least $(m_d + |A| - 3^n)$ points in A.

iii. Let $f : \mathbb{F}_3^n \to \mathbb{F}_3$ such that $f(\mathbf{x}) = 0$ for all $\mathbf{x} \notin A$. Define the matrix $M_f \in \mathbb{F}_3^{A \times A}$ as follows: $M_f(x,y) := f((x+y)/2)$ for all $x, y \in A$. Show that the rank of M_f is exactly $|\{\mathbf{x} \in A | f(\mathbf{x}) \neq 0\}|$.

(c) (4 points) Let $g : \mathbb{F}_3^n \to \mathbb{F}_3$ be a function in $V_d(n)$. Consider the matrix M_g given by $M_g(x,y) := g(x+y)$. Prove the following facts about the rank of the matrix M_g.

i. $\mathrm{rank}(M_g) \leq m_d(n)$.

ii. Strenghten the above to show that $\mathrm{rank}(M_g) \leq 2 \cdot m_{d/2}(n)$.

Hint: Recall that if a $t \times t$-matrix M can be decomposed as $M = UV$ where U is a $t \times r$-matrix and V is a $r \times t$ matrix (or equivalently there exists $2t$ r-dimensional vectors $\mathbf{u}_1, \ldots, \mathbf{u}_t, \mathbf{v}_1, \ldots, \mathbf{v}_t$ such that $M(i,j) = u_i^T v_i$), then $\text{rank}(M) \leq r$.

(d) (1 point) Conclude from the above parts that if A is 3-AP-free, then $|A| \leq m_{2n-d} + 2m_{d/2}$. Setting $d = 4n/3$ show that $|A| \leq 3c^n$ where c is as in Part 5(a)iii