Problem Set 2

- Due Date: **19th October 2025**
- The points for each problem is indicated on the side. The total for this set is **90** points.
- Turn in your problem sets electronically (PDF; either LaTeXed or scanned etc.) via email.
- Collaboration is encouraged, but all writeups must be done individually and must include names of all collaborators.
- Unless explicitly stated (like in question 1 in this set), you are not allowed to use the LLMs for this problem set.
- Referring to sources other than the text book and class notes is strongly discouraged. But if you do use an external source (eg., other text books, lecture notes, or any material available online), ACKNOWLEDGE all your sources (including collaborators) in your writeup. This will not affect your grades. However, not acknowledging will be treated as a serious case of academic dishonesty.
- Be clear in your writing.

Throughout this problem set, we would be working with the inner-product $\langle f, g\rangle = \mathbb{E}[f(v)g(v)]$. All norms, orthogonality etc. are with respect to this inner-product.

1. **[Writing code for zig-zag products]** (15)

   This Google Colaboratory notebook provides a generic implementation of strongly explicit graphs via rotation maps. Create a copy of this notebook, and use Gemini to write code for the tensor product of two graphs, and the zig-zag product of two (compatible) graphs. (Please add your prompts as comments in the code block).

2. **[The Affine Line graph]** (1 + 4 + 2 + 3)

   Let $\mathbb{F}$ be a finite field. Consider the following graph G whose vertex set is $\mathbb{F}^2$ and edges set E defined as
   $$E = \{((a, b), (c, d)) \ : \ ac = b + d\}\,.$$
   One way to interpret this is the point (a, b) is connected to all points (c, d) on the line $y = ax - b$.

   (a) Show that G is $|\mathbb{F}|$-regular.

   (b) Compute the adjacency matrix of the graph G^2. What are its eigenvalues?

   (c) Use the above to show that $\lambda(G) \leq \frac{1}{\sqrt{|\mathbb{F}|}}$.

   (d) Starting with this, and using the graph operations seen in class, show that you can construct a $(D^8, D, 1/8)$-spectral expander for some suitably large constant D.

3. **[Spectral gap for the replacement product]** (3 + 4 + 3)

   In class, we also briefly discussed the *replacement product* of two compatible graphs. If G is a (N, D) graph and H is an (D, d) graph, then the replacement product $G \textcircled{r} H$ is a $(ND, d + 1)$ graph. In this question, we will prove some non-trivial spectral gap bound for $G \textcircled{r} H$ if we know G, H have non-trivial spectral gaps.

Base version: (2025-10-06 19:09:41 +0530) , bb0a0d3

(a) Show that the random walk matrix M of $G \textcircled{r} H$ is a convex combination of Rot_G and $I \otimes M_H$

(b) Write M^3 as a convex combination some operator A and the random walk matrix of $G \textcircled{z} H$, with $\|A\| \leq 1$.

(c) Use the above to argue that $\lambda(G \textcircled{r} H) \geq (p + (1-p) \cdot \lambda(G \textcircled{z} H))^{1/3}$ for some $p \in (0,1)$ that depends only on d.

4. **[Spectral gap for regular, connected, non-bipartite graphs]** (15)

Let G be a D-regular undirected graph on N vertices and let M be the random walk matrix. Suppose $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N \geq -1$ are the eigenvalues of M.

(a) Show that $\lambda_2 = \max_{f \perp \mathbf{1}, \|f\|=1} \langle f, Mf \rangle$. (Note that λ_2 here refers to the second-largest eigenvalue, not necessarily in absolute value).

(b) For any f such that $\|f\| = 1$, prove that $\langle f, Mf \rangle = 1 - \frac{1}{2} \mathbb{E}_{(u,v) \in E}[(f(u) - f(v))^2]$.

(c) Let $f \perp \mathbf{1}$ and $\|f\| = 1$. Suppose G is connected, and f attains its maximum value at a and its minimum value at b, show that

$$1 \leq (f(a) - f(b))^2 \leq \mathrm{poly}(N, D) \cdot (1 - \lambda_2)$$

This immediately provides a lower bound for $1 - \lambda_2$.

[Hint: $f(a) - f(b) = (f(a) - f(u_1)) + (f(u_1) - f(u_2)) + \cdots (f(u_t) - f(b))$ where every term corresponds to some edge.]

(d) Suppose G is a D-regular, connected, non-bipartite undirected graph on N vertices and let $\gamma := 1 - \max(\lambda_2, |\lambda_N|)$. Show that $\gamma \geq \frac{1}{\mathrm{poly}(N,D)}$

[Hint: Consider $G' := G^2$ and use the bounds proved above.]

5. **[Spectrum of complete k-partite graphs]** (5 + 5)

A graph G is said to be a complete k-partite graph if there is a partition of $[N] = V_1 \sqcup \cdots \sqcup V_k$ into k non-empty sets of vertices with the edge set being all possible edges between vertices in two different parts. That is,

$$E = \{(u,v) \ : \ \exists i \neq j \in [k] \ , \ u \in V_i \text{ and } v \in V_j\}$$

(a) Assume that all k parts have exactly N/k vertices each. Show that all non-trivial eigenvalues of the random walk matrix is non-positive.

(b) Prove that even in the general case (where the sizes of V_i need not be the same), all non-trivial eigenvalues of the random walk matrix are non-positive.

6. **[An optimal non-averaging sampler]** (5 + 10)

Suppose $f : \{0,1\}^m \to \{0,1\}$ is some function and $\mu = \mathbb{E}_x[f(x)]$. A (δ, ε)-*sampler* is a randomized algorithm that queries f at various points and outputs some estimate $\hat{\mu}$ with the property that

$$\Pr[|\hat{\mu} - \mu| > \varepsilon] \leq \delta.$$

We are primarily interested in two parameters of such samplers — how many queries did it make, and how many random bits did it use. For this entire problem, assume that we have a *strongly-explicit* $(2^m, d, 0.5)$-spectral expander for some constant d.

(a) Using expanders, show how one can obtain a (δ, ε)-sampler that makes at most $O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\delta}\right)$ queries and uses at most

$$m + O\left(\frac{\log(1/\varepsilon)}{\varepsilon^2} \cdot \log\frac{1}{\delta}\right) \text{ random bits.}$$

(You may assume Theorem 4.22 from Vadhan's manuscript, which is a stronger form of the Question 2 in this set, for this problem. You may also assume that there are strongly explicit constant-degree expanders on 2^m vertices.)

(b) Suppose we have a $((1/8), \varepsilon)$-sampler $\mathcal{S}$ that makes Q queries and uses R random bits. In the last problem set, you studied the "median of averages sampler" built from $\mathcal{S}$:

Run the sampler $\mathcal{S}$ for t independent trials to obtain $\hat{\mu_1}, \ldots, \hat{\mu_t}$. Output the *median* of these estimates.

You should have shown that this new sampler will be an (δ, ε)-sampler if $t = O\left(\log\frac{1}{\delta}\right)$ and instantiated $\mathcal{S}$ using pairwise independence.

Use expander random walks to now construct an (δ, ε)-sampler that makes at most $O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\delta}\right)$ queries and uses at most

$$O\left(m + \log\frac{1}{\varepsilon} + \log\frac{1}{\delta}\right) \text{ random bits.}$$

You have now seen a sampler that has *optimal* number of queries and random bits used (up to constants) but is **NOT** an averaging sampler! Obtaining an averaging sampler with the same performance is an open problem.

7. **Expander walks for bias reduction** (2 + 2 + 8 + 3)

For a random variable X taking values in $\{0,1\}$, define

$$\text{bias}(X) := |\Pr[X=0] - \Pr[X=1]| = \left|\mathbb{E}[(-1)^X]\right|.$$

(2 points) If X is the parity of t independent ε-biased random variables $X_1, \ldots, X_t$ (i.e. $X = X_1 \oplus \cdots \oplus X_t$), show that $\text{bias}(X) \le \varepsilon^t$.

The rest of the problem would be to avoid taking independent XORs and using an expander random walk to reduce bias.

For a function $f : [N] \to \{0,1\}$, let X_f be the associated random variable that takes value 1 with probability $\Pr_v[f(v) = 1]$ and 0 with probability $\Pr_v[f(v) = 0]$. Suppose $\text{bias}(X_f) \le \varepsilon$.

Consider a D-regular, λ-spectral expander G on $[N]$ and let M be the random walk matrix. Define the new random variable X defined by the following sampling process:

Pick a random walk $v_1, \ldots, v_t$ in the graph G. Output $f(v_1) \oplus \cdots \oplus f(v_t)$.

(a) Find an appropriate $N \times N$ matrix F such that you can express $\text{bias}(X)$ as

$$\text{bias}(X) = \left|\langle \mathbf{1}, (FM)^t \mathbf{1}\rangle\right|$$

(b) What upper bound can you give for $\|FM\|$? What about $\|FMFM\|$?

(c) Show that $\text{bias}(X) \le (\varepsilon + \lambda)^{t/2}$.

General note: Take a moment to appreciate the above claim. Parity / XOR is a *really* sensitive function and even a single correlated bit can mess up the parity. It is quite surprising that samples obtained via expander random walks behave like independent samples for parity, even though expander walks are certainly correlated!
