CS 860: Modern Topics in Graph Algorithms	**University of Waterloo: Winter 2024**

Homework 2

Due: Monday, April 8, 2024

Problem 1. Recall the notion of *expander decompositions* from Lecture 6. In this question, we consider expander decompositions for another notion of expansion. For any $\varphi > 0$, we say that a graph $G = (V, E)$ is a φ**-edge-expander** iff

$$\min_{0<|S|\leqslant n/2} \frac{|E(S, \overline{S})|}{|S|} \geqslant \varphi.$$

Notice that φ can be as large as $\Theta(n)$ unlike the expansion measure (conductance) we studied in the class.

Prove that given any $\varphi > 0$, vertices of any graph $G = (V, E)$ can be partitioned into $V = V_1 \sqcup V_2 \sqcup \ldots \sqcup V_k$ such that:

1. For all $i \in [k]$, $G[V_i]$ is a φ-edge-expander.
2. $\sum_{i=1}^{k} |E(V_i, \overline{V_i})| = O(\varphi \cdot n \log n)$.

Notice that unlike in the class, we are bounding the edges by $O(n \log n)$ and not $O(m \log m)$. **(20 points)**

Problem 2. In Lecture 8, we showed that any (β, ε)-EDCS of a *bipartite* graph $G = (V, E)$ contains a $(2/3 - 2\varepsilon)$-approximate matching of G as long as $\beta > 4/\varepsilon$. In this problem, we extend this result to general (not necessarily bipartite) graphs although with some loss in the parameters.

(a) Let $G = (V, E)$ be any arbitrary graph and H be a (β, ε)-EDCS of G for $\beta \geqslant 1000 \cdot \log n/\varepsilon^2$ (feel free to increase the leading constant if needed in your arguments). Let M^* be a maximum matching of G. We create the following partition L, R of V randomly: for every edge $(u, v) \in M^*$ independently, place $u \in L$ and $v \in R$ with probability half and with the remaining probability, place $u \in R$ and $v \in L$. For any vertex not matched by M^*, place it in L or R chosen uniformly at random. Now, let G' be the *bipartite* subgraph of G between vertices L and R and $H' = G' \cap H$.

Prove that with high probability, H' is a (β', ε')-EDCS of G' with $\beta' = (1 + O(\varepsilon)) \cdot \beta$ and $\varepsilon' = O(\varepsilon)$ (for an appropriate choice of constant that you decide). Conclude that H has a $(2/3 - O(\varepsilon))$-approximate maximum matching of G, using the result we proved for bipartite EDCS in the class. **(20 points)**

(b) Use the same randomized process as above but when $\beta \geqslant 1000 \cdot \log(1/\varepsilon)/\varepsilon^2$ (so, we are replacing $\log n$ with $\log(1/\varepsilon)$ only). Use *Lovasz Local Lemma* to prove that with a non-zero probability, H' is a (β', ε')-EDCS of G' with parameters $\beta' = (1 + O(\varepsilon)) \cdot \beta$ and $\varepsilon' = O(\varepsilon)$ (for an appropriate choice of constant that you will decide). Use this to conclude that even in this case (i.e., for a smaller value of β), H has a $(2/3 - O(\varepsilon))$-approximate maximum matching of G. **(20 points)**

Problem 3. Recall the definition of (r, t)-RS bipartite graphs from Lecture 8. In the class, we saw a construction with parameters $r = n/3 - o(n)$ and $t = n^{\Omega(1/\log\log n)}$. In this question, we extend the construction to handle larger induced matchings of size $r = n/2 - o(n)$.

Similar to Lecture 8, we create the vertices of the graph as vectors in $[p]^d$ for $p = d^2$ on each side of the bipartition. Let $\delta < 1/10$ be a parameter.

(*a*) Prove that there exist $t := 2^{\Omega_\delta(d)}$ sets $S_1, \ldots, S_t$ in $[d]$ such that for every $i \neq j \in [t]$, $|S_i| = \delta \cdot d$ and $|S_i \cap S_j| \leqslant 10\delta^2 \cdot d$. **(10 points)**

(*b*) Create an induced matching M_S for each set S of Part (a) similar to the lecture by classifying the vertices in $x \in [p]^d$ based on their "weight" $w_S(x) := \sum_{i \in S} x_i$. However, change the definition of the "bands" such that each *white* band only consists of $|S|$ different classes, while each *red* or *blue* band is considerably larger and consists of $|S|/\delta$ consecutive classes (this part is left vague on purpose for you to figure out exactly how to do this).

Prove that $|M_S| \geqslant n/2 - \delta \cdot n - o(n)$ and all these matchings are induced. **(20 points)**

(*c*) Conclude that for every $\delta \in (0, 1/10)$, there exists an (r,t)-RS graph with $t := n^{\Omega_\delta(1/\log\log n)}$ induced matchings of size $r := (1-\delta) \cdot n/2$ each. **(10 points)**

Problem 4 (**Extra Credit**)**.** Let $G = (V, E)$ be an (r,t)-RS bipartite graph with n vertices on each side of the bipartition. Prove that if $r = n/2$, then t can only be $O(\log n)$. **(+20 points)**

You receive half the credit even for proving a similar statement when $r = (1+\epsilon) \cdot n/2$ for any arbitrarily small constant $\varepsilon > 0$.

Problem 5 (**Extra Credit**)**.** Design an $\widetilde{O}(m)$ time algorithm for finding a (β, ε)-EDCS of any given graph for any *constant* choices of β and ε. **(+20 points)**

Problem 6 (**Extra Credit**)**.** Design an $\widetilde{O}(m \cdot \text{poly}(1/\varepsilon))$ time algorithm for finding a (β, ε)-EDCS of any given graph for any choice of β, say, $\beta = n^{1/4}$ and $\varepsilon > 0$.[1] **(+20 points)**

[1] This question might be too challenging; I certainly do not know the answer to it myself, unless one relaxes the requirements of the EDCS.