Randomized Algorithms 2013A – Problem Set 4

Robert Krauthgamer and Moni Naor

December 31, 2012

1. Recall the Benczur-Karger cut sparsification algorithm seen in class. Suppose the algorithm can use (to determine the probabilities p_e) only an approximation to c_e, say a factor 3 estimate $\hat{c}_e \in [c_e, 3c_e]$. Show how to adapt the algorithm and its analysis.

   Explain the differences without repeating the entire analysis.

2. Recall the Thorup-Zwick distance oracle construction, and prove that for every $v \in V$,

   $$\mathbb{E}[|B(v)|] \leq kn^{1/k}.$$

3. Analyze the following construction for a small data structure that approximates distances within factor 3. Write explcitly the overall storage required (there is no fast query time), and whether the factor 3 in accuracy of queries is worst-case, in expectation, or with high probability.

   Preprocess(G): Choose $L \subseteq V$ as a random set of $l = O(\sqrt{n}\log n)$ "landmark" vertices (for simplicitly, say with repetitions). For every vertex $v \in V$, store its distance (i) to each of the $\sqrt{n}$ vertices closest to it, denoted $B_v \subset V$ (break ties arbitrarily); and (ii) to all the landmark vertices.

   Query(u,v): If $u \in B_v$, i.e., is among the $\sqrt{n}$ closest to v, report the distance. Otherwise, report $\min_{w \in L}[d(u,w) + d(w,v)]$.

   Hint: in the "otherwise" case, try to argue that $L \cap B_v \neq \emptyset$.

Extra credit:

4. Show that for every n there are an n-vertex graph G and some $\varepsilon > 0$, for which every $(1+\varepsilon)$-cut sparsifier G' must have $|E(G')| \geq \Omega(n/\varepsilon)$ (or a similar bound).

   Hint: Consider a complete graph, and start with proving for the case $\varepsilon = 0$ that $|E(G')| \geq \Omega(n^2)$. Then extend it to very small $\varepsilon > 0$.