Randomized Algorithms 2017A – Problem Set 3

Robert Krauthgamer and Moni Naor

Due: Feb. 25, 2017

1. Let $G' = (V, E', w')$ be a $(1+\varepsilon)$-spectral-sparsifier of graph $G = (V, E, w)$ with same $V = [n]$.

   (a) Show that for every dimension $k \geq 1$

   $$\forall x_1, \dots, x_n \in \mathbb{R}^k, \qquad \sum_{uv \in E'} w'_{uv} \|x_u - x_v\|_2^2 \in (1 \pm \varepsilon) \sum_{uv \in E} w_{uv} \|x_u - x_v\|_2^2.$$

   (This generalizes what was seen in class, because the x_i are now k-dimensional vectors.)

   (b) Use the above to prove that for every and $k \geq 2$ and every partition of the vertices $V = V_1 \cup \dots \cup V_k$,

   $$\sum_{i<j} w'(V_i, V_j) \in (1 \pm \varepsilon) \sum_{i<j} w(V_i, V_j).$$

   (This generalizes the cut-sparsification seen in class to k-way cuts.)

2. We saw in class an algorithm for randomized low-diameter decomposition of arbitrary n-point metric space X. Specifically, given a desired diameter bound $\delta > 0$, the probability to separate two points $x, y \in X$ is at most $O(\log n) \cdot \frac{d(x,y)}{\delta}$.

   Show that the same partitioning algorithm achieves the following property (which is obviously stronger): For every $x \in X$ and every radius $\rho \in (0, \delta/8)$,

   $$\Pr_{\Pi \sim \mu} \Big[B(x, \rho) \not\subseteq \Pi(x) \Big] \leq O(\log n) \cdot \frac{\rho}{\delta}.$$

   Note: It suffices to explain the changes to the analysis seen in class.