Sublinear Time Algorithms	March 18, 2019

Homework 3

Lecturer: Ronitt Rubinfeld — *Due Date: April 3, 2019*

Turn in your solution to each *problem on a separate sheet of paper, with your name on each one.*

1. Give a lower bound on computing a multiplicative estimate on the MST of a graph G in adjacency list representation: Give two distributions over graphs of degree at most d and weights in the range $\{1, \ldots, w\}$ (for $w = o(n)$) such that

   (a) graphs in one distribution have an MST weight that is at least twice the MST weight of the graphs in the in other distribution

   (b) in order to distinguish the two distributions with constant probability of success, one must make at least $\Omega(w)$ queries

   If you can get the lower bound to have some nontrivial dependence on d and ϵ, even better!

   (Note: It is possible to write this lower bound without explicitly using Yao's method.) You may assume that G is in adjacency list representation.

2. A *vertex cover* V' of a set of edges E' is a set of nodes such that every edge of E' is adjacent to one of the nodes in V'.

   For graph $G = (V, E)$, let the *transitive closure graph* $TC(G)$ be the graph $G^{(tc)}(V, E^{(tc)})$ where $(u, v) \in E^{(tc)}$ if there is a directed path from u to v in G.

   Let $f : V \to \{0, 1\}$ be a labeling of the vertices of a known directed acyclic graph G by 0 and 1. For any pair of nodes x and y, we say that $x \leq_G y$ if there is a path from x to y in G. We say that f is *monotone* if for all $x \leq_G y$, $f(x) \leq f(y)$. The *minimum distance of f to monotone* is the minimum number of nodes that must be relabeled in order to turn f into a monotone function.

   Let E' be the set of violating edges in $TC(G)$ according to f. Show that the minimum distance of f to monotone is equal to the minimum size of a vertex cover of E'.

3. This problem is about testing monotonicity of functions defined over a directed graph G. The function maps nodes into binary values (i.e., $f : V \to \{0, 1\}$), and we say that it is *monotone* if for all directed edges (u, v), we have that $f(u) \leq f(v)$. We say that f is ϵ-close to monotone if there is a monotone function g such that g and f differ on at most $\epsilon|V|$ entries. A testing algorithm knows the graph G in advance, and for a given node u, may query $f(u)$ in one time step.

   (a) Let $V = \{v_1, \ldots, v_n\}$. For each directed graph $G = (V, E)$, let $B_G = (V', E')$ be the bipartite graph where $V' = \{v_1, \ldots, v_n\} \bigcup \{v'_1, \ldots, v'_n\}$, and $(v_i, v'_j) \in E'$ iff v_j is reachable from v_i in G.

   Show that a q-query testing algorithm for f over graph B_G with distance parameter $\epsilon/2$ yields a q-query testing algorithm for f over graph G with distance parameter ϵ.

(b) Let f be a function on V which is ϵ-far from monotone over graph G. Then $TC(G)$ has a matching of violated edges of size at least $(\epsilon/2)|V|$. (Recall previous problem).

(c) Show that if f is a function over bipartite graph G, there is a test for monotonicity of f with query complexity at most $O(\sqrt{|V|/\epsilon})$.

4. Let $L = \{uu^r vv^r | u, v \in \{0,1\}^*, 2(|u| + |v|) = n\}$. We saw in class that given a string x, distinguishing $x \in L$ from x that is ϵ-far (meaning that $> \epsilon n$ bits of x need to be changed in order to make x a member of L) requires $\Omega(\sqrt{n})$ queries. Give an algorithm for this problem that uses $O(\sqrt{n} \log n / poly(\epsilon))$ queries to the input. The running time does not have to be sublinear.