CMPUT 675 - Assignment #1

Fall 2014, University of Alberta
Due September 24 in class.

Pages: 3

This assignment is to be completed individually. I understand that you may want to discuss the assignment with other students, a good guide for understanding my expectations is that you should not take notes or work out precise details in your discussions (keep it high-level). Mention any discussions and cite any resources you used on the writeup you hand in.

To be clear, whenever you are asked to give an approximation algorithm for a problem, it is expected that you will both describe the algorithm and prove the claimed approximation guarantee. If you can only think of an algorithm with a worse approximation guarantee than I am asking for, then describe it anyway. You may get partial marks (though, it cannot be entirely trivial).

Problem 1)

Marks: 2

In the Maximum Acyclic Subgraph problem, we are given a directed graph $G = (V, E)$. The goal is to find the largest subset F of E such that the graph $H = (V, F)$ contains no directed cycles. Show that there is always a solution F with size at least $|E|/2$ and that such a set can be found in polynomial time. Conclude that this is a $\frac{1}{2}$-approximation.

Problem 2)

Marks: 3

Show that for every $c < \frac{3}{2}$ there is an instance of the Traveling Salesman problem such that Christofides' algorithm will find a solution whose cost is *at least* c times the optimum solution cost. This shows the analysis we saw in class cannot be improved.

Hint:

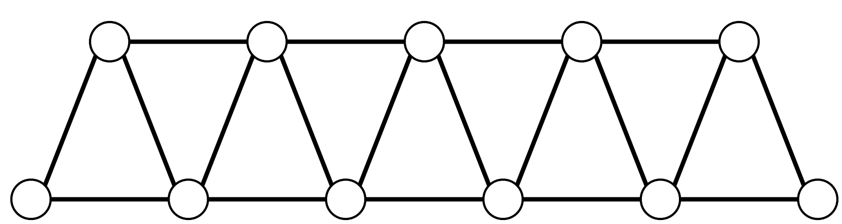

Problem 3)

Marks: 3

Let $G = (V, E)$ be a graph. A *clique* is a subset of nodes $C \subseteq V$ such that $(u, v) \in E$ for every two distinct $u, v \in C$. In the MAXIMUM CLIQUE problem, the goal is to find the largest clique C in a given graph.

Show that if there is an α-approximation for the MAXIMUM CLIQUE problem for some constant $\alpha < 1$, then there is also a $\sqrt{\alpha}$-approximation.

Hint: consider the graph $G^2 = (V \times V, E')$ where $((u, v), (w, x)) \in E'$ for distinct $(u, v), (w, x) \in V \times V$ if and only if both of the following conditions hold:

- $(u, w) \in E$ or $u = w$
- $(v, x) \in E$ or $v = x$.

Side Note: Iterating this argument shows that if there is *some* constant-factor approximation than there is, in fact, a $(1 - \epsilon)$-approximation for *any* constant $\epsilon > 0$. Later in the course, we will see there is, in fact, some constant $c < 1$ such that there is no c-approximation for the MAXIMUM INDEPENDENT SET problem unless P = NP. Considering this lower bound in light of this exercise, we see there is in fact *no* constant-factor approximation unless P = NP.

Problem 4)

Marks: 4

Consider the the k-SUPPLIERS problem, a variant of the k-CENTER problem we discussed in class.

As input, we are given a metric on nodes V with distances $d(u, v), u, v \in V$, a partition of V into two nonempty sets F, C, and an integer $k \geq 1$. We call F the *suppliers* and C the *clients*.

The goal is to find a subset $A \subseteq F$ with $|A| \leq k$ to minimize the maximum distance travelled by a client to their nearest supplier in A. That is, A should minimize $\max\limits_{j \in C} \min\limits_{i \in A} d(i, j)$.

- Give a 3-approximation for the k-SUPPLIERS problem. [3 marks]
- Show that there is no c-approximation for the k-SUPPLIERS problem for any $c < 3$ unless P = NP. [1 marks]

Answer any one of the following two problems. You may attempt both, in which case I will use the highest mark for your grade.

Problem 5)

Marks: 4

Option 1

Consider the ASYMMETRIC TRAVELING SALESMAN problem. The setting is much like classic TSP with the only difference being that the distances are not required to satisfy the symmetry property. That is, we are given a set of locations V and nonnegative costs $c(u, v)$ between $u, v \in V$ that satisfy $c(v, v) = 0$ for $v \in V$ and $c(u, v) \leq c(u, w) + c(w, v)$ for $u, v, w \in V$. However, we may have $c(u, v) \neq c(v, u)$ for locations $u, v \in V$. For example, such distances arise naturally in road networks with one-way streets.

- A *cycle cover* of a subset $U \subseteq V$ is a collection of directed cycles $\mathcal{C}$ such that each $v \in U$ lies on exactly one cycle in $\mathcal{C}$ and each $w \in V - U$ does not lie on any cycle. We allow directed cycles of length two. The cost of $\mathcal{C}$ is the total cost of all edges used by the cycles $\mathcal{C}$.

  Describe how to find a minimum-cost cycle cover for any $U \subseteq V$ with $|U| \geq 2$ in polynomial time. [1 mark]

- Give a $\log_2 n$-approximation for the ASYMMETRIC TRAVLEING SALESMAN problem. [3 marks]

  You may use the following fact without proof. If a directed graph is weakly connected (i.e. the undirected graph obtained by ignoring directions is connected) and for every vertex v the number of edges directed in to v equals the number of edges directed out of v, then there is a tour of the graph that crosses each directed edge exactly once (and in the proper direction).

Option 2

The TRAVELING SALESMAN PATH problem is similar to classic TSP. We are given a (symmetric) metric on locations V with costs $c(u, v)$ between any two $u, v \in V$. Additionally, we are given two distinct nodes $s, t \in V$. The goal is the find the cheapest Hamiltonian path that starts at s and ends at t.

- Let $T = (V, E)$ be a tree and let $D \subseteq V$ be a set of nodes with $|D|$ even. Show that we can efficiently pair the nodes of D into $|D|/2$ pairs $(v_1, v_2), (v_3, v_4), \ldots, (v_{|D|-1}, v_{|D|})$ so that if we let P_i denote the path between the i'th pair, $1 \leq i \leq |D|/2$, then no edge of T lies on more than one of these P_i paths. [1 mark]

- For any spanning tree T of the input where both s and t are leaves of T, let D be the set of nodes in $V - \{s, t\}$ that have odd degree in T. Also, let P be any Hamiltonian $s - t$ path. Finally, let $T + P$ denote the multiset of edges lying on either T or P, where we keep both copies of an edge if it lies on both T and P.

  Show that we can efficiently decompose $T + P$ into three sets of edges E_1, E_2, E_3 such that in each graph $G_i = (V, E_i), 1 \leq i \leq 3$, the set of nodes with odd degree in G_i is precisely D. Such a set of edges is called a *D-join*. [2 marks]

- Give a $\frac{5}{3}$-approximation for the TRAVELING SALESMAN PATH problem. [1 mark]

  You may use, without proof, the fact that if a graph is connected and all but two vertices have even degree, then there is a walk between the two odd-degree nodes that uses each edge exactly once (an *Eulerian walk*). Such a walk can be found in polynomial time.