

MCS 401 – Computer Algorithms I
Fall 2025
Problem Set 6

Lev Reyzin

Due: 12/5/25 by the beginning of class

1. We define the decision version of the Minimum Cut problem as follows: given an undirected, unweighted graph $G = (V, E)$ two distinct vertices $s, t \in V$, and a number k , does there exist an s - t cut of value at most k ?

For each of the two questions below, decide whether the answer is (i) “Yes”, (ii) “No”, or (iii) “Unknown, because it would resolve the question of whether $\mathcal{P} = \mathcal{NP}$.” Give an explanation of your answers.

(a) Question: Is it the case that Minimum Cut \leq_p Vertex Cover?

(b) Question: Is it the case that Independent Set \leq_p Minimum Cut?

2. Suppose someone gives you a black-box B that takes in any undirected graph $G = (V, E)$ and a number k , and in unit time it returns “yes” if G has an independent set of size at least k and “no” otherwise. Design an algorithm with the power to access B as many times as it wishes that returns an independent set of maximum cardinality in a given graph in polynomial time.

3. Consider the problem of making a conference schedule. There are talks T_1, \dots, T_k to be scheduled and participants P_1, \dots, P_ℓ attending the conference. Each participant gives you a list of the talks he is interested in attending. You must schedule times for these talks so that no participant is interested in two talks that are scheduled for the same time. The problem is to determine if a schedule exists that uses only h slots. Show that this problem is NP-complete.

Hint: reduce from graph 3-COLORING, which is defined in the textbook.

4. Consider a set of clauses C_1, \dots, C_k over x_1, \dots, x_n consisting of *unnegated* variables, for example $(x_1 \vee x_3 \vee x_4 \vee x_7)$. Satisfying a collection of clauses over unnegated variables is easy – simply set all the variables appearing in the clauses to 1. However, it is natural to ask whether it is possible to set no more than k variables to 1 and still satisfy such a collection of clauses. Show that this problem is NP-complete.

5. Given objects p_1, \dots, p_n , and distances $d(\cdot, \cdot)$ on them (with $d(p_i, p_i) = 0$, $d(p_i, p_j) = d(p_j, p_i)$, and $d(p_i, p_j) > 0$ for $i \neq j$), the clustering problem of dividing the objects into k sets so as to maximize the minimum distance between any pair of objects in distinct clusters can be solved in polynomial time using minimum spanning trees (see Chapter 4.7 of the textbook).

A different but seemingly related way to formalize the clustering problem would be as follows: divide the objects into k sets so as to minimize the maximum distance between any pair of objects in the same cluster. Whereas the formulation in the previous paragraph sought clusters so that no two were “close,” this new formulation seeks clusters so that none of them is too “wide.”

Given the similarities, it is perhaps surprising that this new formulation is computationally hard to solve optimally. First, let’s write it first as a yes/no decision problem: given n objects p_1, \dots, p_n , with distances on them as before, a number k , and a bound B , we can ask: can the objects be partitioned into k sets, so that no two points in the same set are at distance greater than B from one another?

Prove that this new formulation is NP-complete.

6. Consider the following variant of the 3-SAT problem: given a formula in 3-CNF, the problem is to decide whether there exists an assignment that satisfies exactly all but one of the clauses (and doesn’t satisfy exactly one clause). Show that this problem is NP-complete.