

Exercise Sheet 2

Due date: 14:00, Nov 1st, by the end of the lecture.

Late submissions will be exterminated.

You should try to solve all of the exercises below, and submit two solutions to be graded — each problem is worth 10 points. We encourage you to submit in pairs, but please remember to indicate the author of each individual solution.

Exercise 1 The purpose of this exercise is to show that the greedy algorithm does not always work.

- (a) Explain how you would change Kruskal’s algorithm (as simply as possible) for the Travelling Salesman Problem on complete graphs with positive edge weights: try to greedily build a Hamilton cycle of minimum weight.
- (b) Show that for every $\alpha > 1$, there is an edge-weighted graph for which the greedy algorithm builds a Hamilton cycle whose weight is at least α times larger than the optimum.

Exercise 2 Show that there is a polynomial-time 1.5-approximation algorithm to the Travelling Salesman Problem on non-negatively weighted complete graphs (K_n, ω) satisfying the triangle inequality: for all $u, v, x \in V(K_n)$, $\omega(\{u, v\}) \leq \omega(\{u, x\}) + \omega(\{x, v\})$.

[Hint at <http://discretemath.imp.fu-berlin.de/DMII-2016-17/hints/S02.html>.]

Exercise 3 Assuming $\mathcal{P} \neq \mathcal{NP}$, prove for every $\alpha > 1$ that there is no polynomial-time α -approximation algorithm to the Travelling Salesman Problem on weighted complete graphs.

Bonus (7 pts) Wouldn’t it be nice if the above statement was meaningful? Show that it is by proving $\mathcal{P} \neq \mathcal{NP}$.

Exercise 4 When running Dijkstra's algorithm to find the lightest paths in a weighted graph, we required the edge weights to be non-negative. The following algorithm is designed to find lightest paths in *directed* graphs that may have negative edge weights.

Algorithm: LIGHTPATHS

Data: A directed graph $G = ([n], \vec{E})$, edge weights $\omega : \vec{E} \rightarrow \mathbb{R}$, root vertex $u \in [n]$

Result: LIGHTPATHS(G, ω, u) computes, when possible, for every vertex $v \in [n]$ a lightest directed path (with total weight) from u to v in (G, ω) .

```

/* Initialisation: start with infinite distance bounds and empty
   paths, except for the root u */
Set dist[u] = 0;
for  $v \in [n] \setminus \{u\}$  do
    | Set dist[v] =  $\infty$  ;
    | Set prev[v] = null;
end
/* Repeatedly check edges to see if we can improve current paths */
for  $1 \leq i \leq n - 1$  do
    | for  $edge(x, y) \in \vec{E}$  do // check if edge gives shorter paths from u
        | if  $\text{dist}[x] + \omega((x, y)) < \text{dist}[y]$  then
            | | Set  $\text{dist}[y] = \text{dist}[x] + \omega((x, y))$ ;
            | | Set  $\text{prev}[y] = x$ ;
        | end
    | end
end
/* Run through edges once more to check for %%%%%%%%% %%%% */
for  $edge(x, y) \in E$  do
    | if  $\text{dist}[x] + \omega((x, y)) < \text{dist}[y]$  then
        | | Return error: graph has a %%%%%%%%% %%%%%%%%% ;
    | end
end

```

- (a) Unfortunately the pseudocode got corrupted, and some words were lost. What words should replace the ‘%’ characters towards the end?
- (b) Prove that the algorithm runs correctly. What is its running time?

[Hint at <http://discretemath.imp.fu-berlin.de/DMII-2016-17/hints/S02.html>.]

The next couple of exercises concern SAT — the Boolean satisfiability problem — for which we now define the necessarily terminology. A *Boolean variable* is a variable that can take one of two variables: True or False. A *Boolean formula* $f(x_1, x_2, \dots, x_n)$ is simply a function $f : \{\text{True}, \text{False}\}^n \rightarrow \{\text{True}, \text{False}\}$. In other words, it takes as input a number of Boolean variables, x_1, x_2, \dots, x_n , and for every possible truth assignment to these variables, evaluates to either True or False. A Boolean formula f is said to be *satisfiable* if there is some assignment of truth values to its inputs for which f evaluates to True.

Every Boolean formula can be represented by combining the Boolean input variables with three logical operators: ‘ \wedge ’ (and), ‘ \vee ’ (or) and ‘ \neg ’ (not). In particular, every formula has a *Conjunctive Normal Form (CNF)*. A *literal* is either a variable x_i or its negation $\neg x_i$. A *clause* is the ‘or’ of several literals, so it is satisfied if any one of its literals is. Finally, the CNF formula is the ‘and’ of all its clauses, and is thus satisfied if and only if all of its clauses are. The Boolean satisfiability problem is the decision problem asking whether or not a given CNF formula is satisfiable.

A k -CNF formula consists of clauses with exactly k literals¹, each corresponding to different variables². For example, the following is a 4-CNF formula:

$$f(x_1, x_2, x_3, x_4, x_5, x_6) = (x_1 \vee \neg x_2 \vee x_4 \vee \neg x_5) \wedge (x_2 \vee x_3 \vee \neg x_4 \vee x_6) \wedge (x_1 \vee \neg x_2 \vee x_5 \vee \neg x_6).$$

This formula is satisfiable, since, for example, $f(\text{True}, \text{False}, \text{True}, \text{True}, \text{False}, \text{False}) = \text{True}$. In general, the k -satisfiability (k -SAT) problem is the Boolean satisfiability problem restricted to k -CNFs.

Exercise 5

- (a) Provide an example of an unsatisfiable instance of k -SAT.
- (b) Show that every instance of k -SAT with fewer than 2^k clauses must be satisfiable.

Exercise 6

- (a) Show that 2-SAT is in \mathcal{P} .
- (b) Prove that SAT can be (polynomially) reduced to 3-SAT.

¹Some authors would only ask that there only be at most k literals. However, these are essentially equivalent, since given a clause C with fewer than k literals, we can introduce a new variable y , and replace C with the logically-equivalent $(C \vee y) \wedge (C \vee \neg y)$.

²Having literals with the same variable is redundant: $x \vee x$ is just x , while $x \vee \neg x$ is always True.