

Departement of Computer Science
Johannes Lengler, David Steurer
Lucas Slot, Manuel Wiedmer, Hongjie Chen, Ding Jingqiu

18 December 2023

Algorithms & Data Structures

Exercise sheet 13

HS 23

This exercise sheet is not to be turned in. The solutions will be published at the end of the week, before Christmas.

Exercise 13.1 Shortest path with negative edge weights (part I).


Let $G = (V, E, w)$ be a graph with edge weights $w : E \rightarrow \mathbb{Z}$ and $w_{\min} = \min_{e \in E} w(e)$.

Since Dijkstra's algorithm must not be used whenever some edge weights are negative (i.e., $w_{\min} < 0$), one could come up with the idea of applying a transformation to the edge weight of every edge $e \in E$, namely $w'(e) = w(e) - w_{\min} + 1$, such that all weights become positive, and then find a shortest path P in G by running Dijkstra with these new edge weights w' .

Show that this is not a good idea by providing an example graph G with a weight function w , such that the above approach finds a path P that is not a shortest path in G (this path P can start from the vertex of your choice). The example graph should have exactly 5 nodes and not all weights should be negative.

Exercise 13.2 Shortest path with negative edge weights (part II).

We consider the following graph:



- What is the length of the shortest path from vertex 1 to vertex 6?
- Consider Dijkstra's algorithm (that fails here, because the graph has negative edge weights). Which path length from vertex 1 to vertex 6 is Dijkstra computing? State the sets $S, V \setminus S$ immediately before Dijkstra is making its first error and explain in words what goes wrong.

- (c) Which efficient algorithm can be used to compute a shortest path from vertex 1 to vertex 6 in the given graph? What is the running time of this algorithm in general, expressed in n , the number of vertices, and m , the number of edges?
- (d) On the given graph, execute the algorithm by Floyd and Warshall to find *all* shortest paths. Express all entries of the $(6 \times 6 \times 7)$ -table as 7 tables of size 6×6 . (It is enough to state the path length in the entry without the predecessor vertex.) Mark the entries in the table in which one can see that the graph does not contain a negative cycle.

Exercise 13.3 *Invariant and correctness of algorithm (This exercise is from the January 2020 exam).*

Given is a weighted directed acyclic graph $G = (V, E, w)$, where $V = \{1, \dots, n\}$. The goal is to find the length of the longest path in G .

Let's fix some topological ordering of G and consider the array $\text{top}[1, \dots, n]$ such that $\text{top}[i]$ is a vertex that is on the i -th position in the topological ordering.

Consider the following pseudocode:

Algorithm 1 Find-length-of-longest-path(G, top)

```

 $L[1], \dots, L[n] \leftarrow 0, \dots, 0$ 
for  $i = 1, \dots, n$  do
     $v \leftarrow \text{top}[i]$ 
     $L[v] \leftarrow \max_{(u,v) \in E} \{L[u] + w((u, v))\}$ 
return  $\max_{1 \leq i \leq n} L[i]$ 

```

Here we assume that maximum over the empty set is 0.

Show that the pseudocode above satisfies the following loop invariant $\text{INV}(k)$ for $1 \leq k \leq n$: After k iterations of the for-loop, $L[\text{top}[j]]$ contains the length of the longest path that ends with $\text{top}[j]$ for all $1 \leq j \leq k$.

Specifically, prove the following 3 assertions:

- i) $\text{INV}(1)$ holds.
- ii) If $\text{INV}(k)$ holds, then $\text{INV}(k+1)$ holds (for all $1 \leq k < n$).
- iii) $\text{INV}(n)$ implies that the algorithm correctly computes the length of the longest path.

State the running time of the algorithm described above in Θ -notation in terms of $|V|$ and $|E|$. Justify your answer.

Exercise 13.4 *Cheap flights (This exercise is from the January 2020 exam).*

Suppose that there are n airports in the country Examistan. Between some of them there are direct flights. For each airport there exists at least one direct flight from this airport to some other airport. Totally there are m different direct flights between the airports of Examistan.

For each direct flight you know its cost. The cost of each flight is a strictly positive integer.

You can assume that each airport is represented by its number, i.e. the set of airports is $\{1, \dots, n\}$.

- (a) Model these airports, direct flights and their costs as a directed graph: give a precise description of the vertices, the edges and the weights of the edges of the graph $G = (V, E, w)$ involved (if possible, in words and not formal).

In points (b) and (c) you can assume that the directed graph is represented by a data structure that allows you to traverse the direct predecessors and direct successors of a vertex u in time $O(\deg_-(u))$ and $O(\deg_+(u))$ respectively, where $\deg_-(u)$ is the in-degree of vertex u and $\deg_+(u)$ is the out-degree of vertex u .

- (b) Suppose that you are at the airport S and you want to fill the array d of minimal traveling costs to each airport. That is, for each airport A , $d[A]$ is a minimal cost that you must pay to travel from S to A .

Name the most efficient algorithm that was discussed in lectures which solves the corresponding graph problem. If several such algorithms were described in lectures (with the same running time), it is enough to name one of them. State the running time of this algorithm in Θ -notation in terms of n and m .

- (c) Now you want to know *how many* optimal routes there are to airport T . In other words, if c_{\min} is the minimal cost from S to T then you want to compute *the number of routes from S to T of cost c_{\min}* .

Assume that the array d from (b) is already filled. Provide an as efficient as possible *dynamic programming* algorithm that takes as input the graph G from task (a), the array d from point (b) and the airports S and T , and outputs the number of routes from S to T of minimal cost.

Address the following aspects in your solution and state the running time of your algorithm:

- 1) *Definition of the DP table:* What are the dimensions of the table $DP[\dots]$? What is the meaning of each entry?
- 2) *Computation of an entry:* How can an entry be computed from the values of other entries? Specify the base cases, i.e., the entries that do not depend on others.
- 3) *Calculation order:* In which order can entries be computed so that values needed for each entry have been determined in previous steps?
- 4) *Extracting the solution:* How can the final solution be extracted once the table has been filled?
- 5) *Running time:* What is the running time of your algorithm? Provide it in Θ -notation in terms of n and m , and justify your answer.

Hint: Note that the array d is a part of the input, so you don't need to include the time that is required to fill this array to the running time here.

Exercise 13.5 Elevator.

Consider the following definitions for a **directed** graph $G = (V, E)$:

1. The *out-degree* of a vertex $v \in V$, denoted with $\deg_{\text{out}}(v)$, is the number of edges of E that start at v , i.e., $\deg_{\text{out}}(v) = |\{(v, w) \in E \mid w \in V\}|$.

2. The *in-degree* of a vertex $v \in V$, denoted with $\deg_{\text{in}}(v)$, is the number of edges that end at v , i.e., $\deg_{\text{out}}(v) = |\{(u, v) \in E \mid u \in V\}|$.
3. A *Eulerian walk* is a sequence $v_1, \dots, v_k \in V$ such that $k = |E| + 1$ and $\{(v_i, v_{i+1}) \mid 1 \leq i < k\} = E$. Note that this definition implies (v_i, v_{i+1}) being different edges for $1 \leq i < k$.

In this exercise, you can use without proof the following result:

Lemma 1. A **directed** graph $G = (V, E)$ admits a Eulerian walk if, and only if, all of the following conditions holds:

1. At most one vertex $v \in V$ is such that $\deg_{\text{out}}(v) = \deg_{\text{in}}(v) + 1$;
 2. At most one vertex $v \in V$ is such that $\deg_{\text{in}}(v) = \deg_{\text{out}}(v) + 1$;
 3. Every vertex that satisfies neither (i) nor (ii) is such that $\deg_{\text{out}}(v) = \deg_{\text{in}}(v)$;
 4. The undirected graph G' obtained by ignoring the direction of edges in G is connected.
- (a) Write down the pseudocode of an $O(|V| + |E|)$ time algorithm that takes as input a **directed** graph G , and returns `true` if G has a Eulerian walk, and `false` otherwise. Justify its correctness and complexity.
- (b) Alice is launching iFahrstuhl™, a start-up developing the next generation of elevators. Assume a building with n floors indexed from 1 to n and an elevator which has room for a single person. The elevator receives requests in the form of pairs $(i, j) \in \{1, \dots, n\}^2$ of distinct floors between which a single person is willing to travel.

Consider the scenario where m people want to use the elevator. For $1 \leq t \leq m$, the t -th people want to go from floor i_t to floor j_t . These requests are given as a finite set $S = \{(i_1, j_1), \dots, (i_m, j_m)\}$.

A finite set $S = \{(i_1, j_1), \dots, (i_m, j_m)\}$ of requests is called *optimal* if the pairs can be ordered such that all requests can be processed and the elevator is never empty when moving between two floors (except maybe on its way to fetching the first person).

For example, for $n = 5$, the set $S_1 = \{(2, 3), (4, 1), (3, 4)\}$ is optimal, since it can ordered as $\{(2, 3), (3, 4), (4, 1)\}$, which means that the elevator can start on floor 2 to fetch person 1, go to floor 3, drop person 1 and fetch person 3, go to floor 4, drop person 3 and fetch person 2, go to floor 1, drop person 2, and terminate there. However, the set $S_2 = \{(2, 3), (4, 1)\}$ is not optimal, since there is no way a single elevator can satisfy both requests without moving empty from floor 3 to floor 4 or floor 1 to floor 2.

Given a set of requests S , Alice's elevators should be able to decide whether it's optimal. Model the problem of detecting optimal sets of requests as a graph problem and provide an algorithm to solve it. Describe the vertex and edge set, edge weights (if needed), the graph problem you solve, the algorithm you use, and its complexity. To obtain full points, your algorithm should run in time $O(n + |S|)$.

- (c) Alice's startup has installed k single-person elevators in your n -floor building. Unfortunately, not all elevators can reach all floors. Hence, for each elevator $j \in \{1, \dots, k\}$, you are given a set $F_j \subseteq \{1, \dots, n\}$ of floors it can reach. When you arrive in front of an elevator j , say on floor $f \in F_j$, you can immediately call it, after which you have to wait until it reaches your floor from its current position, moving at the constant speed of 1 time unit per floor. When the elevator arrives, you choose the destination floor $f' \in F_j$, and the elevator brings you to this floor at the constant speed of 0.5 time units per floor (for security reasons, the elevator is slower when it is not empty).

The time spent moving between elevators on the same floor, calling the elevator or choosing the destination floor is negligible, since you are very fast at interacting with elevators.

You are alone in the building at floor 1, with each elevator j being initially located on floor f_j . You would like to go to floor n . What is the minimal amount of time that you have to travel using Alice's elevators? If you cannot reach floor n , then output ∞ .

Model the problem as a graph problem and provide an algorithm to solve it. Describe the vertex and edge set, edge weights (if needed), the graph problem you solve, the algorithm you use, and its complexity. To obtain full points, your algorithm should run in time $O((n + K) \log n)$, where $K = \sum_{j=1}^k |F_j|^2$.


- (d) Continue the setting of (c). Elevator doors in your building need maintenance, but the people in your building also need elevators. In your building, there is exactly one elevator door per elevator and floor, which needs to be functional in order for the elevator to be used from or to this floor. Even if a door is not functional, the elevator can still be used between all other floors where a functional door is present. Alice wants to select as many elevator doors as possible to be maintained during the next working day such that all floors can be reached from each other using the elevators and the remaining functional doors (those not in maintenance).

Model the problem as a graph problem and provide an algorithm to solve it. Describe the vertex and edge set, edge weights (if needed), the graph problem you solve, the algorithm you use, and its complexity. To obtain full points, your algorithm should run in time $O((n + K') \log(n + K'))$, where $K' = \sum_{j=1}^p |F_j|$.

Hint: Consider the set of vertices

$$V = \{v_1, \dots, v_n\} \cup \{w_1, \dots, w_n\} \cup \{\text{elevator}_1, \dots, \text{elevator}_j\}$$

and use subgraphs ("gadgets") of the form



where $F_j = \{i_1, \dots, i_q\}$.