

Problem Set 2

due: October 28, 2019

Exercise 1

15 points

Consider a directed graph $D = (V, A)$, two designated nodes $s, t \in V$, and a cost function $c : A \rightarrow \mathbb{R}_{\geq 0}$.

1. A node potential is a function $\pi : V \rightarrow \mathbb{R}$. For each arc $a = (u, v) \in A$, the reduced costs of a w.r.t. π are defined as $\tilde{c}(a) := \pi(u) + c(a) - \pi(v)$. Show that for two s - t -paths p, q in D holds

$$\tilde{c}(p) \leq \tilde{c}(q) \Leftrightarrow c(p) \leq c(q).$$

2. For a node $v \in V$ let $d(v)$ be the shortest distance from s to v in G . Prove that if D has no negative cycles and all nodes can be reached from s , then the reduced costs w.r.t. d are non-negative.
3. Consider the Shortest Path Instance shown in Figure 1. Run Dijkstra's algorithm to get a shortest path from s to t w.r.t. the costs c . At each iteration specify the label that is extracted from the heap, the labels that are inserted in the heap, and the labels that are updated in the heap. Whenever two labels with the same costs could be extracted from the heap, choose the one that entered the heap later. At the end, specify the found path.




Figure 1: The directed graph $D = (V, A)$.

4. Consider the node potential specified in Table 1. Re-run Dijkstra's algorithm, but in every iteration, instead of picking the label with the current minimum costs, pick the label that minimizes $d(v) + \pi(v)$, were $d(v)$ is the current shortest distance from s to v .

Table 1: Node potential for the SPP instance in Figure 1

v	s	v	u	w	x	y	t
$\pi(v)$	3	4	3	2	2	1	0

Costs transformation and A^* algorithm

The technique showed in Exercise 1.2. shows how to transform an instance's costs, when some of them are negative and no negative cycles exist. After such a transformation, Dijkstra's algorithm can be applied. The potential used for the transformation has to fulfill

$$\pi(u) + c((u, v)) - \pi(v) \geq 0, \forall (u, v) \in A.$$

The node potential given in Table 1 underestimates the shortest path distance from every node to the target node. Such a potential is called *consistent*. As you noticed in Exercise 1.d., running Dijkstra's algorithm w.r.t. the new costs induced by the consistent potential lead to less heap operations. The search is guided towards the target node. This is the main idea behind the A^* algorithm, a widely used variant of Dijkstra's algorithm that often achieves superb running times even for very large graphs.

Exercise 2

5 points

Consider a tree $T = (V, E)$. Show that T cannot have two distinct perfect matchings.