University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
P. Schneider

Advanced Algorithms
Problem Set 10

Issued: Monday, July 15, 2019

Exercise 1: Evaluating Congestion Approximators

As we have seen in the lecture, we can construct a m-congestion approximator based on a maximum spanning tree T as follows ($m := |E|$). For each edge $e \in T$ let S_e be the cut induced by e in the graph. Then we set $R_{e,v} = 1/c_{S_e}$ for all $v \in S_e$ and $R_{e,v} = 0$ for all $v \notin S_e$, where c_{S_e} is the sum of capacities of edges going over the cut S_e. The entries $R_{e,v}$ form a $(n-1) \times n$-matrix R. Show that for $x \in \mathbb{R}^n, y \in \mathbb{R}^{n-1}$ we can compute Rx and $R^\top y$ in $O(n)$. Assume the capacities of the cuts $c_{S_e}, e \in T$ are known.

Exercise 2: Analysis of the Gradient Descent Procedure

In the lecture we saw that we can reduce the max flow problem to a continuous, unrestricted optimization problem that we solved with the gradient descent method.
Show that one step of gradient descent requires $\tilde{O}(m)$ time and one multiplication with R and another one with $R^\top$ (where R is the congestion approximator used in the procedure).

Hint: Use the following chain rule for gradients: for $h(x) := g(Ax)$, we have $\nabla h(x) = A^\top \cdot \nabla g(Ax)$.