Thomas Kesselheim
Anna Heuser

December 8, 2025
Due: December 15, 2025 at 10:00 p.m.

Algorithms and Uncertainty

Winter Term 2025/26

Exercise Set 9

If you would like to present one of your solutions in class, please use the following link to book a presentation slot by Monday evening:
https://terminplaner6.dfn.de/b/1f171f441f362e9dde5926441485b476-1518394
A short meeting to discuss your solution is mandatory before presenting it in class. To book a time slot for this meeting, please use the following link by Monday evening as well:
https://terminplaner6.dfn.de/b/41113685a46a40e79de5bbcb0983c530-1518392

Exercise 1: (3+4+2 Points)

We consider the following modified version of the Boosted Sampling algorithm for Stochastic Steiner Tree from the lecture. The only difference is that it uses ℓ sets $S_1, \ldots, S_\ell$ in the first phase. Show that the approximation guarantee is $\max\{2(1+\frac{\lambda}{\ell+1}), 2(\frac{\ell}{\lambda}+1)\}$. To this end, consider the following tasks concerning the cost of the respective phases.

(a) Give an appropriate analysis for the first phase.

(b) Give an appropriate analysis for the second phase.

(c) Combine both results to derive the desired approximation guarantee.

Exercise 2: (3+4 Points)

In the lecture we show that we can learn for Pandora's Box, by learning the optimal policy for an instance with only slightly shifted costs. We now want to investigate whether it is sufficient instead to learn a good approximation of the fair cap.

a) Prove that obtaining a good approximation of the fair cap implies that the resulting policy is near-optimal. For that consider a policy π' that would be optimal if fair caps were $\sigma'_1, \ldots \sigma'_n$. Show that, if $|\sigma_i - \sigma'_i| \leq \gamma$ for all boxes i, then the expected reward of π' with respect to actual fair caps $\sigma_1, \ldots, \sigma_n$ is at least $V^* - n\gamma$.

b) Now we want to show that the fair cap of an empricial distribution is in general not a good approximation of the true fair cap.

Given a distribution $\mathcal{D}$ of values in $[0,1]$ and cost $c \in [0,1]$, let $\sigma_{\mathcal{D}}$ denote the corresponding fair cap. Additionaly let $\tilde{\sigma}_{\mathcal{D}}$ denote the fair cap with respect to the empirical distribution defined by T samples drawn from $\mathcal{D}$. Show that for any T and any $\epsilon > 0$ there exists a distribution $\mathcal{D}_{(T,\epsilon)}$ and cost $c_{(T,\epsilon)}$ such that $\left|\sigma_{\mathcal{D}_{(T,\epsilon)}} - \tilde{\sigma}_{\mathcal{D}_{(T,\epsilon)}}\right| \geq 1 - \epsilon$ with probability at least $1 - \frac{1}{T}$.