Thomas Kesselheim
Rebecca Lehming

November 18, 2024
Due: November 25, 2024 at 10:00 p.m.

Algorithms and Uncertainty

Winter Term 2024/25

Exercise Set 7

If you want to hand in your solutions for this problem set, please send them via email to rlehming@uni-bonn.de *by Monday evening. Of course, submitting solutions in groups is also possible.*

If you would like to present one of the solutions in class, please also send an email to rlehming@uni-bonn.de *stating* ***which task*** *you would like to present in* ***which of the tutorials****. Deadline for the email is Monday, 10:00 pm. Please note that the tasks will be allocated on a first-come-first-serve basis, so sending this email earlier than Monday evening is highly recommended.*

Exercise 1: (1+4 Points)
Consider the Pandora's Box problem from Lecture 12 but this time we are allowed to keep up to ℓ prizes instead of only one.

a) Define the fair-cap policy for this problem.

b) Show that the fair-cap policy is optimal.

Exercise 2: (3+4 Points)
In order to generalize the Pandora's Box setup from the lecture, suppose we would like to match people $i \in [n]$ to boxes $j \in [m]$ (each person can take at most one prize home). We know that person i's value v_{ij} for the prize in box j is independently drawn from a distribution $\mathcal{D}_{ij}$, but it costs c_{ij} to inspect the exact value of the box v_{ij}. Consider A_{ij}, I_{ij}, σ_{ij}, κ_{ij} and b_{ij} to be the corresponding generalizations of the variables introduced in the lecture.

(a) Show that for any policy π, the expected value is given by

$$V(\pi) = \sum_{i,j} \mathbf{E}\left[A_{ij}\kappa_{ij} - (I_{ij} - A_{ij})b_{ij}\right] \quad .$$

(b) Consider the following generalized policy: Inspect the value of person i for item j v_{ij} in decreasing order of caps σ_{ij}. Every time, the highest observed value so far exceeds the largest remaining cap, i.e.

$$v_{i^*j^*} = \max_{\text{inspected } (i,j)} v_{ij} \quad > \quad \max_{\text{not inspected } (i,j)} \sigma_{ij},$$

we irrevocably match person i^* to box j^* and remove all incident edges to i^* and j^* from the graph.

Show that this policy achieves a value which is at least half the value of the optimal policy.
You can use without a proof that the greedy algorithm for bipartite matching which adds edges in decreasing order of weights to the matching achieves at least half of the optimal max-weighted matching in any bipartite graph.

Exercise 3: (3+4 Points)

We extend the problem of opening boxes from Lecture 13. We are still allowed to open k boxes, but now, we may keep up to ℓ prizes instead of only one.

(a) Derive a linear program such that the expected reward of any adaptive policy is upper-bounded by the value of the optimal LP solution. Give a proof.

(b) Show that the adaptivity gap is still at most 8.