Chair of Operations Research
Department of Mathematics
TU Munich

# Exercise Set 6

Due Wednesday, June 12th, 11:45

Special Topics in Algorithmic Game Theory (MA5226)

Yiannis Giannakopoulos | Diogo Poças | Alexandros Tsigonias-Dimitriadis

Hand in your answers, either in person or by e-mail to `diogo.pocas@tum.de` or `alexandros.tsigonias@tum.de`. **No late submissions accepted.**

**Exercise 6.1** (*Best-response dynamics cycles forever*, Exercise 16.1 from [20LAGT])
Exhibit a game with a pure Nash equilibrium and an initial outcome from which best-response dynamics cycles forever.

**Exercise 6.2** (*Generalized ordinal potential games*, related to Exercises 16.3 and 16.4 from [20LAGT])
A *generalized ordinal potential game* is a cost-minimization game for which there exists a *generalized ordinal potential function*, i.e. a function $\Psi$ such that

$$\text{for any outcome } \mathbf{s}, \text{ agent } i, \text{ and deviation } s_i'\text{: if } C_i(s_i', \mathbf{s}_{-i}) < C_i(\mathbf{s}) \text{ then } \Psi(s_i', \mathbf{s}_{-i}) < \Psi(\mathbf{s}).$$

(a) Extend Proposition 16.1 from [20LAGT] to generalized ordinal potential games. In other words, prove that in generalized ordinal potential games, from an arbitrary initial outcome, best-response dynamics converges to a PNE.

(b) Prove that if best-response dynamics always converges to a PNE (for every choice of initial outcome and beneficial unilateral deviation at each iteration), then the game admits a generalized ordinal potential function.

**Exercise 6.3** (*Equilibrium properties of outcome distribution in no-regret dynamics*, Exercise 17.4 from [20LAGT])
Proposition 17.9 from [20LAGT] proves that the time-averaged joint distribution $\frac{1}{T}\sum_{t=1}^{T}\sigma^t$ generated by no-regret dynamics is an approximate coarse correlated equilibrium, but it says nothing about the outcome distribution $\sigma^t$ in a given iteration $t$. Prove that such a distribution $\sigma^t$ is an approximate coarse correlated equilibrium if and only if it is an approximate (mixed) Nash equilibrium (with the same additive error term).

**Exercise 6.4** (*Best-response dynamics for scheduling games*, Problem 16.1 from [20LAGT])
Recall the class of cost-minimization games introduced in Exercises 4.6 and 5.6 from previous weeks (also Problems 12.3 and 13.1 in [20LAGT]), where each agent $i = 1, 2, \ldots, k$ has a positive weight $w_i$ and chooses one of $m$ identical machines to minimize her load. Consider the following restriction of best-response dynamics:

> **Maximum-Weight Best-Response Dynamics**
>
> While the current outcome $\mathbf{s}$ is not a pure Nash equilibrium:
>
> among all agents with a beneficial deviation, let $i$ denote an agent with the largest weight $w_i$ and $s_i'$ a best response to $\mathbf{s}_{-i}$
>
> update the outcome to $(s_i', \mathbf{s}_{-i})$

Prove that maximum-weight best-response dynamics converges to a pure Nash equilibrium after at most $k$ iterations.

**Exercise 6.5** (*Lower bound on regret*, related to Problem 17.1 from [20LAGT])

Consider an online decision-making problem with $n = 2$ actions. Prove that the worst-case expected regret of an online decision-making algorithm cannot vanish faster than $b/\sqrt{T}$, where $b > 0$ is some constant independent of $T$.

**Exercise 6.6** (*Online decision-making under expert advice*, Problem 17.2 from [20LAGT])

This problem considers a variant of the online decision-making problem. There are $n$ "experts", where $n$ is a power of 2:

**Combining Expert Advice**

At each time step $T = 1, 2, \ldots, T$:

1. each expert offers a prediction of the realization of a binary event (e.g., whether a stock will go up or down)
2. a decision maker picks a probability distribution $p^t$ over the possible realizations 0 and 1 of the event
3. the actual realization $r^t \in \{0, 1\}$ of the event is revealed
4. a 0 or 1 is chosen according to the distribution $p^t$, and a *mistake* occurs whenever it is different from $r^t$.

You are promised that there is at least one omniscient expert who makes a correct prediction at every time step.

(a) A deterministic algorithm always assigns all of the probability mass in $p^t$ to either 0 or 1. Prove that the minimum worst-case number of mistakes that a deterministic algorithm can make is precisely $\log_2 n$.

(b) Prove that for every randomized algorithm, there is a sequence of expert predictions and event realizations such that the expected number of mistakes made by the algorithm is at least $\frac{1}{2} \log_2 n$.

(c) Prove that there is a randomized algorithm such that, for every sequence of expert predictions and event realizations, the expected number of mistakes is at most $b \cdot \log_2 n$, where $b < 1$ is a constant independent of $n$. How small can you take $b$?

This problem set will be discussed in the tutorials on June 14th/19th, 2019.