

Report on the outcomes of a Short-Term Scientific Mission¹

Action number: CA16228

Grantee name: Lucas Baudin

Details of the STSM

Title: No-regret learning in games with queues

Start and end date: 20/02/2022 to 5/03/2022

Description of the work carried out during the STSM

Description of the activities carried out during the STSM. Any deviations from the initial working plan shall also be described in this section.

(max. 500 words)

Prior to the mission, we defined a static model with one queue where the departure time is strategic: players want to send a packet in the queue as late as possible but incur a deterrent penalty if their packets do not reach the end of the queue before the end of the period. Therefore, the sequence of play unfolds during a period of fixed length T . The actions of the players are the departure time (in $[0, T]$). There are I players.

In the beginning of the STSM we refined the results and checked/rewrote proofs I made, under various assumptions, on this static model: coarse correlated (and Nash) equilibria, expectancy on the number of late players. We were especially interested in the case where $I > T$, i.e. there are more players than unit of time in the period, resulting in a game where there is at least one player late. In this case, we proved that there is a single coarse correlated equilibria (and a fortiori a single Nash equilibria) where every player departs on 0 (as early as possible).

In a second part we started studying a dynamic version of this static game (where we only had a general scheme in mind). The dynamic counterpart is a repetition of the static game on many periods but where packets that did not reach the end of the queue are still in the game: players may have several packets and the queue may grow unboundedly (unstable behavior). First, we studied the case where there is *strategic repetition*, that is at every step players use a joint equilibria action profile. In this case, we proved that the queue was stable.

¹ This report is submitted by the grantee to the Action MC for approval and for claiming payment of the awarded grant. The Grant Awarding Coordinator coordinates the evaluation of this report on behalf of the Action MC and instructs the GH for payment of the Grant.

Second, and this was the primary goal of this project, we started investigating what happened if players used no-regret strategies.

Description of the STSM main achievements and planned follow-up activities

Description and assessment of whether the STSM achieved its planned goals and expected outcomes, including specific contribution to Action objective and deliverables, or publications resulting from the STSM. Agreed plans for future follow-up collaborations shall also be described in this section.

(max. 500 words)

In the end of the STSM we decide it was time to write a first draft of a working paper with results on the static model, strategic repetition and no-regret repetition. We think the properties we proved on the static model should be enough to now focus our work on the dynamic model, especially the no-regret repetition. We devised a complete scheme for no-regret learning in this case (main challenges are unbounded payoffs and the varying states) on the blackboard during the STSM. In the days to come, I will finish writing it down.

We currently make strong assumptions on the way no-regret learning is carried out: we suppose that every player uses copies of an exponential-weight algorithm and that these copies are initialized in a specific manner. However, our study of the coarse correlated equilibria (which are usually linked with no-regret learning) of the static game let us think that this should be generalizable. We plan this to be ready to be sent to NeurIPS.