Online Resource Sharing via Dynamic Max-Min Fairness: Efficiency, Robustness and Non-Stationarity

Giannis Fikioris
Cornell University*
gfikioris@cs.cornell.edu

Siddhartha Banerjee
Cornell University
sbanerjee@cornell.edu

Éva Tardos
Cornell University
eva.tardos@cornell.edu

We consider the classical Dynamic Max-min fair (DMMF) mechanism for allocating an indivisible resource without using money over multiple agents and T rounds. We show that under mild assumption on value distributions, the DMMF mechanism guarantees every agent close to optimal utility in large markets.

Our setting is as follows: Each agent i has a fair share α_i (where $\alpha_i > 0$ and $\sum_i \alpha_i = 1$) which represents her priority, e.g., an agent with double the fair share of another agent is twice as important. In every round, each agent has a non-negative value for being allocated the item which is a random variable independent of other agents. The DMMF mechanism allocates as follows: every round a subset of agents requests the resource; out of these, the DMMF mechanism allocates to the one who has received the item for the least number of rounds, normalized by their fair shares.

An agent with a fair share of α_i can hope to win the α_i fraction of the rounds that correspond to her highest values. Using this idea, we define an agent i's *ideal utility* $v_i^\star$ as the maximum expected average per-round utility she gets if there were no other agents and she can win only an α_i fraction of the rounds. We use $v_i^\star T$ as a benchmark for how much utility agent i can hope to get over T rounds since this is the maximum utility she can gain without violating the other agents' fair shares.

We show that in the DMMF mechanism, any agent i can use a simple strategy to make utility guarantees that hold under arbitrary (potentially adversarial and collusive) behavior by the other agents. Specifically, when her values are i.i.d. across rounds, agent i can guarantee total expected utility by every round $t \leq T$:

- $\frac{1}{2} v_i^\star t - O(1)$ under any value distribution.
- $(1 - \sqrt{2\alpha_i}) v_i^\star t - O(1)$ when her values are uniformly distributed.
- $(1 - \sqrt{2\lambda\alpha_i}) v_i^\star t - O(1)$ when her value distribution's pdf max and min value over its support differ by a $\lambda \geq 1$ multiplicative factor.

The last two results are approximately optimal in the large market setting, where $\alpha_i \to 0$ and therefore agent i can guarantee almost $v_i^\star T$ utility.

We also prove that an agent can have utility guarantees under the same mechanism, even when her values are dependent across rounds, generated by a hidden Markov model. We use a parameter $\gamma \in (0, 1]$ to measure the dependence of values across rounds, where $\gamma = 1$ indicates that values are independent and lower γ indicates more dependence across rounds. In this case, we prove that agent i can always guarantee $\Omega(\gamma) v_i^\star t - O(1)$ utility for small α. We also offer guarantees that are independent of γ.

Our results improve over previous work, which only provides a $\frac{1}{2} v_i^\star T - O(\sqrt{T})$ utility guarantee for the last round T for the i.i.d. setting, and using a more complex mechanism. In addition to having stronger guarantees, our mechanism can handle dynamic populations by resetting the mechanism when new agents arrive or leave. If the mechanism is reset m times then the additive error is at most $O(m)$ even if these resets happen arbitrarily, since our utility guarantees hold in every round and have constant additive loss.

We also study the reusable resource setting, where an agent might need the item for multiple consecutive rounds to receive any utility. With minor modifications to the DMMF mechanism, we can make the same utility guarantees as the i.i.d. setting, albeit only for the last round and with a $O(\sqrt{T})$ additive error.

*Part of this work was done while Giannis Fikioris was visiting Google as a Student Researcher in the Market Algorithms team.