Indo/US Collaborative Research Grants

National Science Foundation of US and Technology Innovation Hubs of India

सत्यमेव जयते

Title: No-regret Algorithms for Fair Resource Allocation

Indian PI: Prof. Abhishek Sinha, School of Technology and Computer Science, Tata Institute of Fundamental Research.

US PIs: Prof. Mohammad Hajiesmaili, Manning College of Information and Computer Sciences University of Massachusetts Amherst.

In the **online fair allocation problem,** users sequentially request resources from a fixed pool of resources shared among the users over multiple rounds. The objective is to learn how to allocate resources best so that the most number of requests are satisfied over all the rounds. Problems in various domains like **online caching**, **job scheduling,** and **matching** can be modeled in this way. However, only optimizing for total requests served by any policy leads to a biased allocation of resources among users where some users are served much more requests than others. In this project, we propose an online learning policy that guarantees a fair allocation of resources among users while still having high efficiency. To do so, we optimize the ***α-fair utility function*** which induces a trade-off between the desired efficiency and fairness by incorporating a notion of diminishing return property in the global objective function. Previously, it was shown that for any online policy, the difference between the aggregate α-fair utilities of the agents between an optimal static clairvoyant allocation (which has access to the full request sequence in advance) and that of the online policy grows at least linearly with time. That is, no policy can achieve sublinear regret in this setting. **Our proposed policy is the first policy that achieves c-approximate sublinear regret with the approximation factor c ≤ 1.445**. The upper bound on the c-regret exhibits a surprising phase-transition phenomenon at α=0.5. The proof of our results **introduces new algorithmic and analytical techniques, including greedy estimation of the future gradients for non-additive global reward functions and bootstrapping adaptive regret bounds, which may be of independent interest**. We validate the performance of our algorithm by performing numerical experiments on online caching and scheduling using synthetic and real-world datasets. Our algorithm achieves good overall utility while being fair in terms of the popular *Jain's fairness index* compared to other algorithms.