

Batching and Optimal Multi-stage Bipartite Allocations

Thursday, September 26, 2024 3:30 PM (45 minutes)

In several applications of real-time matching of demand to supply in online marketplaces — for example, matching delivery requests to dispatching centers in Amazon or allocating video ads to users on YouTube — the platform allows for some latency (or there is an inherent allowance for latency) in order to batch the demand and improve the efficiency of the resulting matching in such non-stationary environments. Motivated by these scenarios, I investigate the optimal trade-off between batching and inefficiency in the context of designing online allocations under non-stationary arrivals in this talk. In particular, I consider a variant of the classic adversarial online bipartite allocation family of problems where demand requests arrive stage-wise and in K batches—in contrast to online arrival. Our main result for this family is an optimal $(1 - (1 - 1/K)^K)$ -competitive (fractional) matching algorithm, improving the classic $(1 - 1/e)$ competitive ratio bound known for the online variants of these problems (Mehta et al., 2007; Aggarwal et al., 2011; Golrezaei et al. 2014). Time permits, I briefly explain a generalization of the base result to multistage configuration allocations and its applications to video display advertising and AdX.

Our main technique at a high level is developing algorithmic tools to vary the trade-off between “greediness” and “hedging” of the matching algorithm across stages. We rely on a particular family of convex-programming-based matchings that distribute the demand in a specifically balanced way among supply in different stages, while carefully adjusting the balancedness of the resulting matching across stages. More precisely, we identify a (unique) sequence of polynomials with decreasing degrees to be used as strictly concave regularizers of the basic linear program for the greedy matching at each stage. At each stage, our multi-stage algorithm returns the corresponding regularized optimal solution as the matching of this stage (by solving the convex program). By providing structural decomposition (tied to these convex programs) of the underlying subgraph at each stage and recursively connecting the regularizers together, we develop a new multi-stage primal-dual framework to analyze the competitive ratio of this algorithm. For the extension to multi-stage configuration allocation, we develop a new idea of regularizing separately at different price levels, and we analyze the resulting regularized price-level convex programming-based multistage algorithm using a new primal-dual argument.

Presenter: NIAZADEH, Rad (University of Chicago)

Session Classification: Afternoon Session