Sequential Fair Allocation of Limited Resources under Stochastic Demands

Sean R. Sinclair[*], **Gauri Jain**, **Siddhartha Banerjee** and **Christina Lee Yu**
School of Operations Research and Information Engineering, Cornell University
{srs429, gj82, sbanerjee, cleeyu}@cornell.edu

Introduction

Our work here is motivated by a problem faced by our local food-bank (Food Bank for the Southern Tier of New York (FBST)) in operating their mobile food pantry program. Every day, FBST uses a truck to deliver food supplies directly to distribution sites (soup kitchens/pantries/etc.). When the truck arrives at a site, the operator observes the demand there and chooses how much to allocate before moving to the next site. The number of people assembling at each site changes from day to day, and the operator typically does not know the demand of later sites (but has a sense of the demand distribution based on previous visits). Finally, the amount of food in the truck is usually insufficient to meet the total demand, and so the operator must under-allocate at each site, while trying to be *fair* across all sites. The question is:

What is a fair allocation here, and how can it be computed?

In *offline* problems, where demands (more generally, utility functions) for all agents are known to the principal, there are many well-studied notions of fair allocation of limited resources. A relevant notion in our context is that a fair allocation is one satisfying two desiderata: *pareto-efficiency* (for any agent to benefit, another must be hurt) and *envy-freeness* (no agent prefers an allocation received by another). This definition draws its importance from the fact that in many allocation settings, it is both known to be achievable, and also to encompass other natural desiderata (in particular, *proportionality*, wherein each agent's utility is at least that achieved under equal allocation). In particular, when goods are divisible, then for a large class of utility functions, an allocation satisfying both is easily computed (via a convex optimization program) by maximizing the Nash Social Welfare (NSW) objective subject to allocation constraints.

Many settings, much like the FBST operating their mobile food pantry, have principals make decisions *online*, with incomplete knowledge on the demands for agents to come. However, these principals have access to historical data allowing them to generate demand histograms for each agent. Designing allocation algorithms in this setting necessitates utilizing the Bayesian information of the demand distribution to ensure equitable access to the resource, while adapting to the online realization of demands as it unfolds. Guaranteeing pareto-efficiency and envy-freeness simultaneously is impossible in this setting. However, it is important to develop algorithms which achieve probabilistic version of fairness by utilizing the distributional knowledge to develop algorithms that are *approximately fair*.

Overview of our Contributions

We first formalize the online stochastic fair allocation problem described above, and demonstrate that in the online setting, there are distributions for which no policy can achieve pareto-efficiency and envy-freeness over all realizations. This motivates studying approximate notions.

For any allocation, a measure of its (un)fairness is the maximum of the additive deviation of the realized utilities in terms of envy-freeness and (normalized) pareto-efficiency. Further, in the online setting, a natural fair-allocation policy is one which minimizes the expected value (over all realizations) of this score for the ex-post allocation. Minimizing this objective can be formulated as a Markov decision process (MDP); however, the resulting policy may be computationally expensive to find, and also, difficult to interpret.

In order to develop practical heuristics for the above objective, we consider an alternate objective of minimizing $\mathbb{E}[\|X^{opt} - X^{alg}\|_\infty]$, the expected ℓ_∞ deviation in hindsight between the allocation X^{alg} made by the algorithm, and the offline NSW maximizing solution X^{opt}. We show that an ϵ approximation for this objective gives a $c\epsilon$-approximation for the fairness score defined above (where c is a problem dependent constant). The usefulness of this is not immediately apparent, as minimizing this objective may harder given it has a higher dimensional state-space. However, we show how this gives rise to natural heuristic policies based on solving *information-relaxed* optimization problems, where relevant random quantities are replaced with their expectation. These policies are simple, scalable, and in experiments, generate allocations close to the optimal fair allocation in hindsight (and in particular, outperforms the optimal online policy for maximizing NSW). Moreover, it also does well across time, with the minimal difference in performance between earlier and later arriving agents. Thus, we believe this policy is a promising candidate for practical online fair allocation.

---

[*]Contact Author

This is a short one-page abstract for the workshop, for more information please email the authors.