Contribution ID: **35** Type: **not specified**

Stable Matching: Choosing Which Proposals to Make

Tuesday, July 11, 2023 4:45 PM (20 minutes)

Ishan Agarwal and Richard Cole

Abstract: To guarantee all agents are matched, the classic Deferred Acceptance algorithm needs complete preference lists. In practice, preference lists are short, yet stable matching still works well. This raises two questions:

Why does it work well?
Which proposals should agents include in their preference lists?

We study these questions in a model, introduced by Lee, with preferences based on correlated cardinal utilities: these utilities are based on common public ratings of each agent together with individual private adjustments. Lee showed that for suitable utility functions, in large markets, with high probability, for most agents, all stable matchings yield similar valued utilities. By means of a new analysis, we strengthen Lee's result, showing that in large markets, with high probability, for all but the agents with the lowest public ratings, all stable matchings yield similar valued utilities. We can then deduce that for all but the agents with the lowest public ratings, each agent has an easily identified length O(log n) preference list that includes all of its stable matches, addressing the second question above. We note that this identification uses an initial communication phase.

We extend these results to settings where the two sides have unequal numbers of agents, to many-to-one settings, e.g. employers and workers, and we also show the existence of an epsilon-Bayes-Nash equilibrium in which every agent makes relatively few proposals. These results all rely on a new technique for sidestepping the conditioning between the tentative matching events that occur over the course of a run of the Deferred Acceptance algorithm. We complement these theoretical results with an experimental study.

Presenter: AGARWAL, Ishan

Session Classification: Track A-1