On the Social Welfare of Mechanisms for Repeated Batch Matching

(Extended Abstract)

Elliot Anshelevich
Rensselaer Polytechnic Institute
eanshel@cs.rpi.edu

Meenal Chhabra
Rensselaer Polytechnic Institute
chhabm@cs.rpi.edu

Sanmay Das
Rensselaer Polytechnic Institute
sanmay@cs.rpi.edu

Matthew Gerrior
Rensselaer Polytechnic Institute
gerrim@cs.rpi.edu

ABSTRACT

We study hybrid online-batch matching problems, where agents arrive continuously, but are only matched in periodic rounds, when many of them can be considered simultaneously. Agents not getting matched in a given round remain in the market for the next round. This setting models several scenarios of interest, including many job markets as well as kidney exchange mechanisms. We consider the social utility of two commonly used mechanisms for such markets: one that aims for stability in each round (greedy), and one that attempts to maximize social utility in each round (max-weight). Surprisingly, we find that in the long term, the social utility of the greedy mechanism can be higher than that of the max-weight mechanism. We hypothesize that this is because the greedy mechanism behaves similarly to a soft threshold mechanism, where all connections below a certain threshold are rejected by the participants in favor of waiting until the next round. Motivated by this observation, we propose a method to approximately calculate the optimal threshold for an individual agent, based on characteristics of the other agents, and demonstrate empirically that social utility is high when all agents use this strategy.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics

Keywords

Matching, Search, Social welfare

1. INTRODUCTION

Many matching scenarios operate in a hybrid online/batch mode, where agents arrive and wait until the next market clearing period. In any given clearing period, all candidates currently waiting are considered for a match. Those who are successfully matched leave the market, while others wait for the next clearing period. This describes scenarios ranging from kidney exchange (which clear every few weeks) to academic job markets (typically once a year) [2, 3, 1].

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

2. MODEL AND STRATEGIES

Time proceeds in discrete steps, and at each unit of time (a *round*) all agents who are thus far unmatched participate in a batch matching. At time $t = 0$ there are n agents, and at each future time period r new agents arrive. The agents connect to each other with probability p. Agents can be thought of as nodes on a graph. The existence of an edge between two nodes means there is a non-zero utility to both from being matched with each other. An edge between two agent i and j is associated with a weight u_{ij} that determines the utility of that matching. It is assumed that agents lose utility by a factor of δ ($\delta \in (0\ 1)$) per time unit for waiting.

Social utility is additive, and given by:

$$U = \sum_{i,j \in \text{Matches}} u_{ij}(\delta^{t-t_i} + \delta^{t-t_j})$$

where t_i and t_j are the arrival times of agent i and agent j respectively, and t is the time at which they are matched.

We assume that u_{ij}'s are i.i.d draws from a stationary distribution $f(x)$ irrespective of the type of the agents to be connected and the time at which the edge is formed.

At each round, all unmatched agents can report their set of acceptable neighbors to the mechanism, so the mechanism finds an acceptable matching. Once this reporting is done, the agents cannot change their mind and have to accept the match chosen by the mechanism. Unmatched agents are eligible to be matched again in the next round. From an agent's perspective, selecting the acceptable matches can be seen as a sequential search problem. We can show that, under certain conditions, an agent's optimal strategy is the same in any round, and can be characterized by a reservation value t^* such that the agent should (pre-)reject all potential matches with utility less than t^* and be ready to accept any match with utility greater than t^*. The optimal threshold t^*

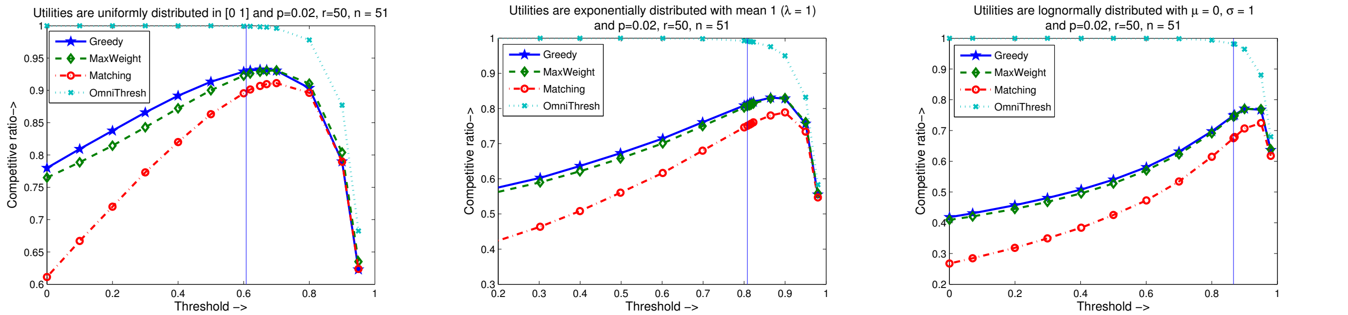

U1:Uniform($U[0\,1]$, $p = 0.02$) E1:Exponential ($\lambda = 1$, $p = 0.02$) L1:Lognormal(($\mu = 0, \sigma = 1$), $p = 0.02$)

Figure 1: Competitive ratio of social welfare (compared with the Omniscient matching) as a function of threshold. Surprisingly, greedy matching yields higher social welfare than max weight matching at lower thresholds. All the curves are unimodal which shows that there exists only one optimal threshold.

can be calculated by solving the following Bellman equation:

$$t^* = \delta(t^* \Pr(\neg \mathrm{M}) + \mathbb{E}(\mathrm{Utility}|\mathrm{M}) \Pr(\mathrm{M})) \qquad (1)$$

where M represents the event that the agent is matched with another with the utility of the match being greater than t^*. Quantifying $\Pr(\mathrm{M})$ is difficult. Therefore, we propose an algorithm to approximately calculate t^*. In Figure 1, the vertical blue line represents the threshold t^* calculated using our algorithm.

We calculate the social utility using the following set of mechanisms:

Online Maximum Weight Matching: The matching at each round is formed using the max-weight matching algorithm, using only edges such that $u_{ij} > \tau_{ij}$

Online Greedy: Similar to Online Maximum Weight Matching, except that matchings are formed using the greedy algorithm in each round. This is also roundwise stable.

Online Maximal Matching: After removing all edges below τ_{ij}, the mechanism picks an arbitrary maximal matching.

Omniscient Matching: This mechanism has foresight into the future and calculates the optimal solution to the offline problem.

OmniThresh Matching: This mechanism gives us an upper bound on the overall social utility of any threshold-based offline algorithm.

3. EXPERIMENTAL RESULTS

Figure 1 demonstrates the empirical performance of the threshold based mechanisms and the proposed threshold calculation algorithm. We observe the following:

1. Threshold mechanisms can significantly improve social welfare. Figure 1 shows the improvement in social welfare due to using threshold mechanisms. The vertical blue line represents the approximately optimal threshold t^* discussed in Section 2, and is close to the best threshold for maximizing social welfare (as well as to the best threshold for rational agents to use, as shown above).

In Figure 1, we also observe the unimodal behavior of the competitive ratio w.r.t threshold. This indicates that the social welfare exhibited by our thresholded online mechanisms is most likely a combination of two effects which counteract each other: (1) Having a high enough threshold removes some of the "online" nature of the mechanism, since it no longer matches pairs on low-quality edges, and instead waits to match them in future rounds, and (2) Having a high enough threshold removes high-quality edges from consideration, thus making a matching worse.

2. Greedy performs better than Max-Weight. Another interesting property apparent from the Figure 1 is the fact that the Greedy mechanism consistently performs better than the Max-Weight mechanism. The Greedy mechanism guarantees stability, while the Max-Weight mechanism maximizes social welfare. Thus, it seems surprising that, in aggregate, the Greedy mechanism is superior.

3. Thresholds matter more than edge weights: support for unweighted matching. Figure 1 shows that, while the Online Maximal Matching mechanism performs worse than the mechanisms that take actual edge weights into account, it still performs well (often within just a few percent of the other online mechanisms) when the threshold is picked appropriately.

Finally, we note that in order to scale to more realistic domains like kidney matching, our model needs to accommodate agents of multiple types. Our initial experiments with two agent types (the types are characterized by their probabilities of connecting to agents of all types) are promising: they suggest that threshold mechanisms, with thresholds chosen appropriately for each type, may continue to work well with multiple types. For further discussion and analysis of our results, we direct the reader to the full version of our paper.

4. ACKNOWLEDGEMENTS

The authors acknowledge support from NSF grants IIS-0952918, CCF-1101495 and CCF-0914782, and US-Israel BSF grant no. 2008404.

5. REFERENCES

[1] A. Abdulkadiroglu, P. Pathak, and A. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *Amer. Econ. Rev.*, 99(5):1954–1978, 2009.

[2] D. Abraham, A. Levavi, D. Manlove, and G. O'Malley. The stable roommates problem with globally-ranked pairs. *Intern. and Net. Econ.*, pages 431–444, 2007.

[3] A. Roth and E. Peranson. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, 1999.