

Preference Elicitation in Matching Markets via Interviews: A Study of Offline Benchmarks

(Extended Abstract)

Baharak Rastegari
School of Computing Science
University of Glasgow
Glasgow, UK
baharak.rastegari@glasgow.ac.uk

Paul Goldberg
Dept. of Computer Science
University of Oxford
Oxford, UK
paul.goldberg@cs.ox.ac.uk

David Manlove
School of Computing Science
University of Glasgow
Glasgow, UK
david.manlove@glasgow.ac.uk

ABSTRACT

In this paper we study *two-sided matching markets* in which the participants do not fully know their preferences and need to go through some costly deliberation process in order to learn their preferences. We assume that such deliberations are carried out via *interviews*, thus the problem is to find a good strategy for interviews to be carried out in order to minimize their use, whilst leading to a stable matching. One way to evaluate the performance of an interview strategy is to compare it against a naïve algorithm that conducts all interviews. We argue however that a more meaningful comparison would be against an optimal offline algorithm that has access to agents' preference orderings under complete information. We show that, unless $P=NP$, no offline algorithm can compute the optimal interview strategy in polynomial time. If we are additionally aiming for a particular stable matching, we provide restricted settings under which efficient optimal offline algorithms exist.

Keywords

Two-sided matching; preferences; interviews

1. INTRODUCTION

Two-sided matching markets model many practical settings, such as corporate hiring and university admission [6, 3]. In the classical *stable marriage problem* participants are partitioned into two disjoint sets, and each participant on one side of the market wishes to be matched to a candidate from the other side of the market and has preferences over potential matches. A matching is called stable if no pair of participants would prefer to leave their assigned partners to pair with each other. Gale and Shapley's seminal paper [1] proposed a polynomial-time algorithm for finding a stable matching; a rich literature has developed since.

A key assumption in much of this literature is that all market participants know their full preference orderings. However, as markets grow large (e.g., in the hospital-resident matching market or college admission market [1, 3]) it quickly becomes impractical for participants to assess their precise

preference rankings. Instead, participants usually start out with some partial knowledge about their preferences and need to perform some deliberation in order to learn their precise preference ordering. In this paper we assume that deliberations are carried out via *interviews*. Interviews are usually costly, therefore we wish to minimize their usage.

Any interviewing strategy leads to refinements of the partial orders contained in the original problem instance that represented uncertainty over the true preferences. A key aim could be to carry out sufficient interviews so as to arrive at an instance that admits a super-stable matching μ . Informally speaking, super-stability ensures that μ will be stable regardless of how the remaining uncertainty is resolved. The original instance need not admit a super-stable matching (see [2] for an example) but we are guaranteed that a super-stable matching is always achievable (e.g., by conducting all possible interviews). We seek a good strategy that conducts as few interviews as possible so as to obtain a refined instance that admits a super-stable matching. In general any such strategy will be an online algorithm, since the next interview to be carried out might depend on the results of previous ones.

2. PRELIMINARY DEFINITIONS

SMP, SMTI and SMT.

In an instance of the *Stable Marriage problem with Partially ordered preferences (SMP)*, there are two sets of agents, namely a set of men M and a set of women W . Each agent a finds a subset of agents on the opposite side of the market acceptable – we refer to these as a 's *acceptable candidates*. An agent a 's preferences over his/her acceptable candidates need not be strict. That is, given two candidates, a might be indifferent between them. We denote by I an instance of SMP, and by p_a the partial orders that represent the preference ordering of agent a . A *matching* μ is a pairing of men and women such that each man is paired with at most one woman and vice versa, and no agent is matched to an unacceptable partner. A well studied special case of SMP is the *Stable Marriage problem with Ties and Incomplete lists (SMTI)*. In SMTI, each agent has a partition of acceptable candidates into *ties* such that he or she is indifferent between the candidates in the same tie but has a strict preference ordering over the ties. The *Stable Marriage problem with Ties (SMT)* is the special case of SMTI in which each man finds each woman acceptable and vice versa.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Interviews to refine the partial orders.

In a given instance I of SMP in this paper, we assume that the partial preference ordering profile represents the agents' initial information state. That is, agents may not have enough information initially in order to rank their acceptable candidates in strict order. However each agent a has a strict preference ordering \succ_a over his or her acceptable candidates, although s/he may not initially be aware of this entire ordering. We let $\succ_{M,W}$ denote the *strict (true underlying) preference ordering profile of all agents*. The task of the agents is to learn enough information about their acceptable candidates in order to refine their preferences, in a manner consistent with $\succ_{M,W}$, to obtain an SMP instance I' that admits a super-stable matching μ (thus μ will be stable with respect to $\succ_{M,W}$).

Following the model introduced in [4], we assume that instances can be refined through *interviews*. Each interview pairs one man m with one woman w . An interview is informative to both parties involved. When agent a interviews ℓ candidates, this results in a new refined SMP instance which is exactly the same as I except that a now has a strict preference ordering over the ℓ interviewed candidates.

Note that not all refinements of I can be reached by a set of interviews. We say that an SMP instance I' is an *interview-compatible refinement* of I if I' can be refined from I using interviews. We define the *cost* of I' to be the minimum number of interviews required to refine I into I' .

Definitions of the interview minimization problems.

The motivating problem is as follows: given an instance I of SMP, find an interview-compatible refinement I' of minimum cost such that I' admits a super-stable matching. Since the result of one interview might influence which interview/s to carry out next, any strategy for carrying out interviews should be regarded as an online algorithm. Towards computing bounds for the competitive ratio of an online algorithm, the offline scenario is of interest, and that is what we consider in what follows. In the offline case, the mechanism designer is given $\succ_{M,W}$, the strict (true underlying) preference ordering profile of the agents, and would like to compute an optimal interviewing schedule, i.e., an interview-compatible refinement I' of I , such that $\succ_{M,W}$ refines I' . This is reflected in the definition of the following problem, named MIN-ICR, which is an abbreviation for "Minimum-cost Interview Compatible Refinement problem".

DEFINITION 1. *An instance of MIN-ICR comprises a tuple $(I, \succ_{M,W})$, where I is an instance of SMP and $\succ_{M,W}$ is a strict preference ordering profile that refines I . The problem is to find an interview-compatible refinement I' of I such that (i) $\succ_{M,W}$ refines I' , (ii) I' admits a super-stable matching, and (iii) I' is of minimum cost amongst interview-compatible refinements that satisfy (i) and (ii).*

It is sometimes the case that we aim for a particular matching, stable under $\succ_{M,W}$, that has some desirable properties, for example it is preferred to every other stable matching by women. The offline problem can then be viewed as a restricted variant of MIN-ICR where, in addition to I and $\succ_{M,W}$, we are also equipped with a matching μ . This is reflected in the definition of the following problem, named MIN-ICR-EXACT, which is an abbreviation for "Minimum-cost Interview Compatible Refinement problem with Exact matching".

DEFINITION 2. *An instance of MIN-ICR-EXACT comprises a tuple $(I, \succ_{M,W}, \mu)$, where I is an instance of SMP, $\succ_{M,W}$ is a strict preference ordering profile that refines I , and μ is a matching that is weakly stable w.r.t. $\succ_{M,W}$. The problem is to find an interview-compatible refinement I' of I , such that (i) $\succ_{M,W}$ refines I' , (ii) μ is super-stable in I' , and (iii) I' is of minimum cost amongst interview-compatible refinements of I that satisfy (i) and (ii).*

3. RESULTS

In the full version of this paper [5], we prove that MIN-ICR and MIN-ICR-EXACT are NP-hard even if I is an instance of SMTI in which each tie is of size at most 3. Further, we prove that both problems are NP-hard even for SMT instances, and even if all men are indifferent between all women. The hardness proofs are by reduction from the Vertex Cover problem (VC).

We then explore the tractability of MIN-ICR-EXACT under various restricted settings, using a reverse reduction from MIN-ICR-EXACT to VC to prove our claims. We show that MIN-ICR-EXACT is solvable in polynomial time under three different restrictions: (i) if one side has fully known strict preference ordering, (ii) if I is an instance of SMTI in which ties are of size at most 2, and (iii) if I is an instance of SMTI in which all men are endowed with the same ties, as well as all women.

4. FUTURE WORK

The main direction is to investigate the online case, where the true underlying preferences are not known to the mechanism designer, with respect to measures such as the competitive ratio. Furthermore, an important question is whether MIN-ICR is polynomial-time solvable under some restricted setting. Extending the known results on interviewing in stable marriage markets to many-to-one markets such as college admission is another important future direction. It is also interesting to study online algorithms in a setting where elicitation is taking place via comparison queries.

Acknowledgments. We thank Rob Irving and Piotr Krysta for the useful discussions and their valuable feedback. This work was supported by EPSRC grants EP/K01000X/1 and EP/K010042/1, and COST Action IC1205.

REFERENCES

- [1] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Amer. Math. Monthly*, 69:9–15, 1962.
- [2] R. W. Irving. Stable marriage and indifference. *Discrete Applied Math.*, 48(3):261–272, 1994.
- [3] D. F. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific, 2013.
- [4] B. Rastegari, A. Condon, N. Immorlica, and K. Leyton-Brown. Two-sided matching with partial information. In *Proc. EC '13*, pp. 733–750. ACM, 2013.
- [5] B. Rastegari, P. Goldberg and D. Manlove. Preference elicitation in matching markets via interviews: A study of offline benchmarks. CoRR abs/1602.04792, 2016.
- [6] A. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.