Mechanisms for the Marriage and the Assignment Game*

Paul Dütting[1] and Monika Henzinger[2]

[1] Ecole Polytechnique Fédérale de Lausanne (EPFL), Station 14, CH-1015 Lausanne
`paul.duetting@epfl.ch`
[2] Universität Wien, Fakultät für Informatik, Universitätstrasse 10/11, A-1090 Wien
`monika.henzinger@univie.ac.at`

Abstract. Starting with two models fifty years ago, the discrete *marriage game* [1] and the continuous *assignment game* [2], the study of *stable matchings* has evolved into a rich theory with applications in many areas. Most notably, it has lead to a number of truthful mechanisms that have seen a recent rejuvenation in the context of sponsored search. In this paper we survey the history of these problems and provide several links to ongoing research in the field.

Keywords: stable matchings, auction mechanisms, matching markets

1 Introduction

Starting with two models fifty years ago, the discrete *marriage game* [1] and the continuous *assignment game* [2], the study of *stable matchings* has evolved into a rich theory with applications in many areas. Most notably, it has lead to a number of truthful mechanisms that have seen a recent rejuvenation in the context of sponsored search. In this paper we survey the history of these problems and provide several links to ongoing research in the field.

Both models are given as input two sets, I and J, and have to *match* or *assign* the elements of I to the elements of J. The goal is to find a *stable matching* between the elements of I and the elements of J. The exact definition of what constitutes a stable matching depends on the model under consideration. For many models the next question to study is which stable matching is the "best" stable matching, using a notion of optimality that has a useful economic interpretation. This leads then to the question of *truthfulness*, i.e., whether players can achieve even "better" outcomes for themselves if they misreport the model parameters related to them.

The literature on stable matchings can be classified as follows:

1. One-to-one, one-to-many, or many-to-many: To how many elements of set J can an element of set I be matched simultaneously, and vice versa? We say

---

* This work was conducted as part of a EURYI scheme award (see http://www.esf.org/euryi/).

that a matching is *one-to-one* if the elements of both sets can be matched with at most one element of the other set. It is *one-to-many* if the elements of one set can be matched to more than one element of the other set, but not vice versa. If elements of both sets can be matched to multiple elements of the other set, then it is *many-to-many*.

2. Preference lists or utility functions: Are the preferences that each element of I resp. J has with respect to the elements of the other set given as (not necessarily complete) preference lists, or are they given as utility functions?
3. One-sided or two-sided: Is the notion of stability *one-sided*, i.e., the only concern is to respect the preferences of one of the two sets, or is it *two-sided*, i.e., do we we have to respect the preferences of both sets?
4. Further restrictions: If we think of one of the two sets as "the buyers" and of the other as "the sellers" and assume that each seller has a unique item to sell, are there any restrictions on the prices of the items? Do the sellers impose *reserve prices*, i.e., lower bounds on the price they are willing to accept? Or do the buyers have a *maximum price* or *budget*, i.e., an upper bound on the price they are willing to pay?

In this survey we restrict ourselves to discussing one-to-one matchings. We first define all game-theoretic terminology in Section 2. Afterwards, in Section 3 and 4, we describe the results for one-to-one matchings in chronological order. We conclude with open research problems in the area.

2 Game-Theoretic Terminology

In this section we introduce the game-theoretic terminology used in this survey.

A *game* is defined by a set N of *players*, each with a set S_i of possible strategies. In our case the set of players N is given by the sets I and J. The set of strategies S_i of player i is given by the preferences that she can specify, depending on the model either as a preference list or as a set of utility functions.

Each game has an *outcome* o. In our case the outcome specifies a matching μ, i.e., it tells us which elements are matched, and - in the case where the preferences are given as utility functions - a price p for each pair (i, j) in the matching. Given such an outcome o, the players in N derive a certain *utility*. In the case of preference lists, these utilities are given by the corresponding entries in the preference lists. In the case of utility functions, the utilities are given by the corresponding utility functions.

A strategy profile s specifies for each player i a strategy s_i. Such a strategy profile s^* is a (weakly) *dominant strategy* [3] if for all i, all $s_i \in S_i$, and all $s_{-i} \in S_{-i} : u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$, where $u(s_i^*, s_{-i})$ resp. $u(s_i, s_{-i})$ denotes the utility of player i if she plays s_i^* resp. s_i and the other players play $s_{-i} \in S_{-i} = S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \ldots S_{|N|}$. That is, a strategy profile is a dominant strategy if every player is guaranteed to get the highest possible utility, no matter which strategy the other players choose. If in a mechanism truthful revelation of preferences is a dominant strategy, then the mechanism is said to be *incentive compatible* [3], or *truthful* [4]. We will use these terms interchangeably.

The *core* [3] of a game is defined in terms of a *value function* v that associates with each subset of players P a value $v(P)$. It is defined as the set of feasible payoff vectors that cannot be improved upon. Formally: A vector $(x_i)_{i \in P}$ of real numbers is a *P-feasible payoff vector* if $\sum_{i \in P} x_i = v(P)$. We refer to a N-feasible payoff vector as *feasible payoff vector.* A feasible payoff vector $(x_i)_{i \in N}$ *cannot be improved upon*, if there is no coalition P and P-feasible payoff vector $(y_i)_{i \in P}$ for which $y_i > x_i$ for all $i \in P$.

3 Marriage Game

The model with discrete preferences is typically referred to as *marriage game.* It was introduced by Gale and Shapley in 1962 [1]. The goal is to match n men with m women in a one-to-one fashion. The preferences are given as (not necessarily complete) preference lists. A matching is *stable* if no two men α and β are assigned to women A and B, although both would prefer to be matched to the other partner, i.e., man α prefers woman B over A and woman A prefers man β to man α. Note that the stability notion is *two-sided* as both, men and women, discriminate between the partners of the opposite sex. A stable matching is *men-optimal* (or *women-optimal*) if every man (or woman) is at least as well off under it as under any other stable matching. A simple example shows that a matching that is optimal for both, men and women, does not exist in all cases.

The main contribution of Gale and Shapley was a combinatorial mechanism, known as the *deferred acceptance mechanism*, that always finds a men-optimal (or women-optimal) matching. The mechanism works in rounds. In each round all men (or women) who have not yet been accepted propose to their favorite women (or men). Each woman keeps her favorite from among those who have proposed to her, but does not accept him yet. She does not accept him to allow for the possibility that someone better may propose to her later. She rejects all other men. In the next round all men who have been rejected propose to their next best choices, and so on. The running time of this mechanism is polynomial.

Dubbins and Freedman [5] and later Gale and Sotomayor [6] showed that if an mechanism computes a men-optimal (or women-optimal) matching, then it is a *dominant strategy* [3] for the men (or women) to reveal their true preferences. In other words: It is in the best interest of the men (or women) to reveal their true preferences, because they will end up with a partner that they like at least as much as the partner they would get if they misreported their preferences. That is, any such mechanism is *truthful.*

There exists many variants of the marriage game that we do not want to discuss in detail here (for a detailed discussion see [7]), some of which do not permit a polynomial time mechanism unless $P = NP$ [8].

4 Assignment Game

In the *assignment game*, which was introduced by Shapley and Shubik in 1972 [2], n buyers are to be matched with m sellers. Each seller has a unique item to sell, so the terms “seller” and ”item” can be used interchangeably. The goal is to compute a matching μ along with a vector of prices p. In the *two-sided* version both, the buyers and the sellers, derive a certain utility from being matched to a seller or buyer at a given price, where the utility of the buyers (resp. sellers) is a monotonically decreasing (resp. increasing) function of the price. In the *one-sided* version the seller does not discriminate between buyers. Hence the one-sided version is a special case of the two-sided version, where the utility function of each of the sellers is the same for all buyers. The utility functions are typically assumed to be continuous. Depending on which form the utility functions have we distinguish the following cases.

4.1 Linear Case (With Slope 1). This is the case that was originally studied by Shapley and Shubik. The i-th buyer values the item of the j-th seller at $v_{i,j}$ dollars. The j-th seller values her item at r_j dollars. The value r_j can be interpreted as the *reserve price* imposed by the seller, because she will not sell her item below this price. If seller j sells her item to buyer i for p_j dollars, then j's final profit (or utility) is p_j and i's profit (or utility) is exactly $u_i = v_{i,j} - p_j$. A matching μ with prices p is *stable* if (1) $u_i + p_j = v_{i,j}$ if i and j are matched, and (2) $u_i + p_j \geq v_{i,j}$ if they are not. Since the sellers do not discriminate between the buyers, this model can be characterized as being one-sided.

Shapley and Shubik formulated the problem as a linear program and used this formulation to (1) prove that a stable outcome always exists and (2) identify the set of stable outcomes with the *core* [3] of the associated game. In this specific case the value $v(P)$ of a subset of players P is $v(P) = \max[v_{i_1,j_1} + \cdots + v_{i_k,j_k}]$, where $k = \min(|I \cap P|, |J \cap P|)$ and the maximum is taken over all arrangements of $2k$ distinct players $i_1, \ldots, i_k$ in $P \cap I$ and $j_1, \ldots, j_k$ in $P \cap J$. A payoff (u', p') is P-feasible if $\sum_{i \in P} u'_i + \sum_{j \in P} p'_j = v(P)$. The core is the set of N-feasible payoffs (u, p) for which there is no P-feasible payoff (u', p') in which a all buyers and sellers in P are strictly better off, i.e., $u'_i > u_i$ for all $i \in I \cap P$ and $p'_j > p_j$ for all $j \in J \cap P$. So intuitively, if you think of the game as a bipartite graph with node sets I and J and edges $(i, j) \in I \times J$ with weight $v_{i,j}$, then this shows that every outcome in the core must be a maximum weight matching.

They also showed that the set of stable outcomes has a *lattice structure*, i.e., for any two payoffs (u', p') and (u'', p'') that correspond to stable outcomes the payoffs $(\underline{u}, \overline{p})$ and $(\overline{u}, \underline{p})$ with $\underline{u_i} = \min(u'_i, u''_i)$, $\overline{p_j} = \max(p'_j, p''_j)$, $\overline{u_i} = \max(u'_i, u''_i)$, and $\underline{p_j} = \min(p'_j, p''_j)$ also correspond to stable outcomes. That is, any two stable outcomes can be combined in a way that is “best” for all buyers resp. “best” for all sellers. This “lattice lemma” allows them to prove the existence of highest payoffs u^*_i (resp. p^*_j) and lowest payoffs u_{*i} (resp. p_{*j}) in the core. These payoffs correspond to the buyer- and seller-optimal outcomes.

A similar approach, which is also based on LPs, was taken by Leonard in 1983 [9], who showed how to “tweak” the LPs so that by solving first the primal program and then the dual program constrained to the outcome of the primal program one can find the stable outcome with the *highest* (*lowest*) payoff to the buyers. He also showed that any mechanism that computes this outcome is *truthful* for the buyers (sellers).

The first combinatorial mechanism for this problem is due to Demange et al. [10]. [3] This mechanism, which they refer to as *multi-item auction*, is based on the Hungarian Method of Kuhn [11, 12]. It starts with an initial price vector p_0 in which all prices are set to the reserve prices r. Each bidder i now announces which item(s) j he wants to buy at the current prices. Call each such pair (i, j) a *desired edge*. If it is possible to create a one-to-one matching consisting only of desired edges, this matching is already the desired stable outcome. If no such assignment exists then the mechanism raises the price of some of the items. By a celebrated theorem from combinatorics, namely Hall's Theorem [13], a necessary and sufficient condition for a stable outcome is that there is no overdemanded item set. An item set S is *overdemanded* if the number of bidders all of whose desired edges go to items in S is greater than the number of items in S. It is *minimal* if none of its proper subsets is overdemanded. By raising the prices of all items in a minimal overdemanded item set by the same amount the mechanism either (a) resolves the overdemand for this item set or (b) increases the size of the minimal overdemanded item set under consideration. Since the size of this set is bounded by the number of items, this shows that eventually all overdemand will be resolved and, hence, a stable outcome is found.

In fact, this mechanism does not only find a stable outcome, but the stable outcome with the "smallest" prices, which corresponds to the buyer-optimal outcome. It is truthful, because in this case *any* mechanism that finds a buyer-optimal outcome is truthful [9]. The running time can be shown to be polynomial if in each step the prices are raised "by as much as possible".

Recently, variants of the "classic" assignment game have been studied, in which (1) each seller has a *reserve price* $r_{i,j}$ for each of the buyers and (2) each buyer has a *maximum price* or *budget* $m_{i,j}$ for each of the items. In [14] it was shown that if the input is in "general position", then a unique buyer-optimal outcome exists and can be found in polynomial time using a variant of the Hungarian Method. This mechanism is *truthful* for inputs in "general position". In [15] it was shown that, with a slightly different notion of stability, the "general position" requirement can be dropped, i.e., a buyer-optimal outcome *always* exists. The authors of [15] also showed how to compute this outcome in polynomial time using a different variant of the Hungarian Method. This mechanism – just as the mechanism of [14] – is *not* truthful for arbitrary inputs.

4.2 Piece-Wise Linear Case. A more general case, namely that of "locally linear" or *piece-wise linear* utilities has been studied by Alkan in 1989 [16] using an LP-based approach and in 1991 [17] using a combinatorial approach.

A piece-wise linear utility function is a (continuous) function that is composed of linear segments. Each segment can have a different slope. With utilities of this form the above observation that by raising the prices of all items in a minimal

---

[3] In their paper Demange et al. argue that their model is "unsymmetrical" in that each buyer specifies number of items many numbers (her valuations), while each seller specifies only one number (the reserve price). They claim that there mechanism can be modified to handle buyer-dependent reserve prices. This would make the model two-sided, as now the sellers would discriminate between the buyers.

overdemanded item set by the same amount it possible to either (a) resolve the overdemand or (b) increase the size of the minimal overdemanded set is no longer true. The main contribution of [16] was to prove the existence of a "stable direction" or amounts by which the prices have to be raised in the piece-wise linear case so that this observation continues to hold.

The main contribution of [17] was to use these "stable directions" for the price updates in the *multi-item auction* of [10]. This mechanism finds the buyer-optimal outcome and runs in time polynomial in the number of buyers and items, and linear in the number of "linear domains" entered. A linear domain is a region of the price space in which the slope of all piece-wise linear utility functions remains the same.

Note that this result also establishes the existence of a buyer-optimal stable outcome for arbitrary, but continuous utility functions (discussed in more detail below), because any such function can be approximated arbitrarily well by piece-wise linear utility functions. In [18] it is shown that any mechanism that finds a buyer-optimal outcome for arbitrary, but continuous utility functions is *truthful*. Since piece-wise linear utilities are a special case of general utilities this shows that Alkan's generalization of the *multi-item auction* is truthful.

4.3 General Case. The case of arbitrary, but continuous utility functions was studied by Demange and Gale [18] in 1985. They were able to generalize the "lattice lemma" of [2] by showing that even in this very general case (and with a two-sided notion of stability) the set of stable outcomes forms a *lattice*, which by a similar argument as in the linear case, establishes the existence of a buyer- resp. seller-optimal outcome. Using the buyer- resp. seller-optimality and a lemma that was originally proved by Hwang, they proved that the buyer- resp. seller-optimal outcome cannot be manipulated by a subset of the buyers resp. sellers. That is, for all buyers it is a dominant strategy to reveal their true preferences. In other words: Any mechanism that computes a buyer- resp. seller-optimal outcome is *truthful* for the buyers resp. sellers. Demange and Gale did not give a mechanism to compute such outcomes.

The variant of this problem with not necessarily continuous utility functions, that are, e.g., required to model *maximum prices* or *budgets*, has been addressed only recently [19], where it was shown that the *lattice structure* persists in the presence of discontinuities and implies the existence of a buyer-optimal outcome. The authors of [19] also gave a brute-force mechanism that finds this outcome in time exponential in the number of items, but polynomial in the number of buyers.

5 Open questions

There are several open questions: First of all, it would be interesting to see whether Alkan's approach for continuous piece-wise linear utilities can be extended to settings with discontinuous piece-wise linear utilities. Furthermore, it would be interesting to (dis)prove the existence of a polynomial time mechanism for computing a buyer-optimal outcome for general (continuous or discontinuous) utility functions. Finally, it would be interesting to study similar questions for the one-to-many and many-to-many case that we did not discuss here.

References

1. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. The American Mathematical Monthly **69**(1) (1962) 9–15
2. Shapley, L., Shubik, M.: The assignment game I: The core. International Journal of Game Theory **1**(2) (1972) 111–130
3. Myerson, R.B.: Game Theory: Analysis of Conflict. Harvard University Press (1997)
4. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press (2007)
5. Dubins, L.E., Freedman, D.A.: Ms. Machiavelli and the Gale-Shapley algorithm. American Mathematical Monthly **88**(7) (1981) 485–494
6. Demange, G., Sotomayor, M.A.O.: Some remarks on the stable matching problem. Discrete Applied Mathematics **11**(3) (1985) 223–232
7. Roth, A.E., Sotomayor, M.A.O.: Two-Sided Matching - A Study in Game-Theoretic Modelling and Analysis. Cambridge University Press (1990)
8. Manlove, D.F., Irving, R.W., Iwama, K., Miyazaki, S., Morita, Y.: Hard variants of stable marriage. Theoretical Computer Science **276**(1–2) (2002) 261–279
9. Leonard, H.B.: Elicitation of honest preferences for the assignment of individuals to positions. The Journal of Political Economy **91**(3) (1983) 461–479
10. Demange, G., Gale, D., Sotomayor, M.A.O.: Multi-item auctions. The Journal of Political Economy **94**(4) (1986) 863–872
11. Kuhn, H.: The Hungarian Method for the assignment problem. Naval Research Logistics Quarterly **2**(1-2) (1955) 83–97
12. Frank, A.: On Kuhn's Hungarian Method - A tribute from Hungary. Naval Research Logistics **51**(1) (2004) 2–5
13. Hall, P.: On representatives of subsets. J. London Math. Soc. **10**(2) (1935) 26–30
14. Aggarwal, G., Muthukrishna, S., Pál, D., Pál, M.: General auction mechanism for search advertising. World Wide Web Conference (2009) 241–250
15. Dütting, P., Henzinger, M., Weber, I.: Sponsored search, market equilibria, and the Hungarian Method. Symposium on Theoretical Aspects of Computer Science (2010) 287–298
16. Alkan, A.: Existence and computation of matching equilibria. European Journal of Political Economy **5**(2-3) (1989) 285–296
17. Alkan, A.: Equilibrium in a Matching Market with General Preferences. In: Equilibrium and Dynamics. Essays in Honour of David Gale. The MacMillan Press Ltd (1992)
18. Demange, G., Gale, D.: The strategy structue of two-sided matching markets. Econometrica **53**(4) (1985) 873–883
19. Dütting, P., Henzinger, M., Weber, I.: Bidder optimal assignments for general utilities. Workshop on Internet and Network Economics (2009) 575–582