Available online at www.sciencedirect.com
ScienceDirect
Theoretical Computer Science 379 (2007) 317–328
Theoretical Computer Science
www.elsevier.com/locate/tcs

Polynomial time algorithm for an optimal stable assignment with multiple partners

Vipul Bansal[a,*], Aseem Agrawal[b], Varun S. Malhotra[c]

[a] *Adobe Systems, I-1A Sector 25A, Noida 201301, India*
[b] *Citigroup Global Markets, 3 Temasek Avenue, Singapore 039190, Singapore*
[c] *Oracle Corporation, Redwood Shores, CA 94065, USA*

Abstract

This paper considers the many-to-many version of the stable marriage problem where each man and woman has a strict preference ordering on the members of the opposite sex that he or she considers acceptable. Further, each man and woman wishes to be matched to as many acceptable partners as possible, up to his or her specified quota. In this setup, a polynomial time algorithm for finding a stable matching that minimizes the sum of partner ranks across all men and women is provided. It is argued that this sum can be used as an optimality criterion for minimizing total dissatisfaction if the preferences over partner-combinations satisfy a no-complementarities condition. The results in this paper extend those already known for the one-to-one version of the problem.

Keywords: Stable marriage; Multiple partners; Optimal assignment; Substitutes; Dis-satisfaction score; meta-rotation

1. Introduction

The stable assignment problem, first described by Gale and Shapley [4] as the stable marriage problem, involves an equal number of men and women each seeking one partner of the opposite sex. Each person ranks all members of the opposite sex in strict order of preference. A matching is defined to be stable if no man and woman prefers each other to a current partner. Gale and Shapley showed the existence of at least one stable matching for any instance of the problem by giving an algorithm for finding it. An introductory discussion of the problem is given by Polya et al. [7] and an elaborate study is presented by Knuth [11]. Variants of this problem have been studied by Gusfield and Irving [6] amongst others, including cases where the orderings are over partial lists or contain ties. It is known that a stable matching can be found in each of these cases individually in polynomial time (Gale and Sotomayor [5], Gusfield and Irving [6]). However in the case of simultaneous occurrence of incomplete lists and ties, the problem becomes NP-hard (Iwama et al. [10]).

While the search version of the problem is shown to be polynomially solvable in most situations, the problem of counting the number of possible stable matchings for a given problem instance is exponential. A useful concept

* Corresponding author.
E-mail addresses: vbansal@adobe.com (V. Bansal), aseem.agrawal@citigroup.com (A. Agrawal), varun.malhotra@oracle.com (V.S. Malhotra).

0304-3975/$ - see front matter
doi:10.1016/j.tcs.2007.02.050

relevant in this context is that of #P-completeness (see Valiant [20,21]). Irving and Leather [8] showed that the corresponding enumeration problem is #P-complete.

McVitie and Wilson [13] pointed out that the algorithm by Gale and Shapley [4], in which men propose to women, generates a *male-optimal* solution in which every man gets the best partner he can in any possible stable matching and every woman gets the worst partner she can in any stable matching. They suggested an egalitarian measure of optimality under which the sum of the ranks of partners for all men and women was to be minimized. Irving et al. [9] provided an efficient algorithm to find a stable matching satisfying the optimality criterion of McVitie and Wilson [13].

The present work extends the work of Irving and Leather [8] and Irving et al. [9] to the many-to-many version of the problem, where each man or woman may have multiple partners. Such a situation may arise in the context of matching hospitals with doctors (consultants), buyers with sellers and similar other cases.

In a general many-to-many matching problem, a person may have preferences defined over subsets of the members of the opposite set. The literature on many-to-many and one-to-many stable matching problems can be grouped into two categories based on the assumptions placed on the preference function of a person over the members of the opposite set.

One approach, coming from the economics domain, assumes that each person (or firm) specifies a strict preference ordering defined on all possible subsets of the set of acceptable partners (or workers). Further, workers and firms may regard each other as *substitutes*, that is, if a worker is a desirable employee to a firm amongst a subset of workers, then he/she continues to be so even amongst a less desirable subset of workers. Under these two assumptions, many results developed for the one-to-one stable matching have been shown to have their counterparts for one-to-many (Roth and Sotomayor [18]) and many-to-many situations (Roth [17], Sotomayor [19], Martinez et al. [12] and Alkan [2,1]). The substitutable preferences model, however, has a computational limitation due to exponential nature of the preference function which puts a lower bound on what any algorithm can achieve.

The other approach comes from the computer science domain and is closely related to the original *stable marriage* formulation of the problem by Gale and Shapley [4]. Here, each man or woman has an upper limit on the number of partners that he or she may be assigned and specifies a preference ordering only on the *individuals* of the opposite sex that he or she considers acceptable and not on all possible combinations of them. This approach is simpler, computationally attractive and well suited for situations where it is feasible to rank only the individual items. Taking this approach, Baiou and Balinski [3] showed the existence of male and female optimal assignments under strict preference orderings and provided characterizations for them. However, they did not address the issues relating to complexity of enumeration of the stable marriages. Furthermore, determining the stable matchings under an equitable measure of optimality remains an open problem.

In this paper, we work with the second approach to the many-to-many stable matching problem and generalize the notion of optimality proposed for one-to-one matching by McVitie and Wilson [13]. We show that the optimality criterion makes sense provided that we include a no-complementarities condition for preferences on combinations of partners. Although this may seem to resemble the first approach, it differs significantly in not requiring the specification of preference functions of exponential size and makes do with only preference ordering on individuals of the opposite set. We then generalize the methodology described by Irving and Leather [8] and Irving et al. [9] and show the existence of the corresponding results for the many-to-many stable marriage problem. In particular, we obtain a polynomial time algorithm for finding an optimal assignment and show how all the stable matchings for the problem can be enumerated. The algorithm and the other results are generic and are not dependent upon the no-complementarities assumption. In addition to the results themselves, the paper also reveals a novel methodology (which we call *meta-rotations*, extending the concept of *rotations* by Irving and Leather [8]), which is a generic technique with potential use in solving search problems.

In the next section, we formally describe our model of the multiple partner stable marriage problem. Subsequently, we introduce our optimality criterion and provide the conditions and justification for its usefulness. We then propose a methodology for search space reduction and show that it leads to a polynomial time algorithm for finding an 'optimal' matching. Finally, some concluding remarks are presented.

2. Multiple partner stable marriage model

Let $M = \{m_1, \ldots, m_{|M|}\}$ and $F = \{f_1, \ldots, f_{|F|}\}$ respectively denote the sets of $|M|$ males and $|F|$ females. Every person has a strict preference order over those members of the opposite sex that he or she considers acceptable. Let L_m

be the preference list of male m with the most preferred partner being the first in the list. Similarly, L_f represents the preference ordering of female f. Incomplete lists are allowed so that $|L_m| \le |F|$, $\forall m \in M$ and $|L_f| \le |M|$, $\forall f \in F$. Each person also has a quota on the total number of partners with which he or she may be matched. Furthermore, a person prefers to be matched to a person in its preference list than to be matched to fewer than the number of persons specified by his or her quota. Let q_m and q_f denote the respective quotas of male m and female f. The many to many stable marriage problem can thus be specified by $P = (M, F, L_M, L_F, Q_M, Q_F)$ where L_M and L_F respectively denote the $|M| \times 1$ and $|N| \times 1$ vectors of male and female preference lists, and Q_M and Q_F represent the $|M| \times 1$ and $|N| \times 1$ vectors of male and female quotas respectively.

Example 1. Consider an instance P specified by: $|M| = 6, |N| = 6$

$$\begin{array}{llll}
L_{m_1} = [f_2, f_6, f_1, f_3, f_5, f_4] & L_{f_1} = [m_5, m_1, m_2, m_3, m_4, m_6] \\
L_{m_2} = [f_1, f_2, f_4, f_5, f_6, f_3] & L_{f_2} = [m_1, m_2, m_3, m_4, m_5] \\
L_{m_3} = [f_5, f_3, f_1, f_4, f_2, f_6] & L_{f_3} = [m_3, m_4, m_5, m_1, m_2, m_6] \\
L_{m_4} = [f_4, f_1, f_2, f_3, f_6, f_5] & L_{f_4} = [m_6, m_3, m_4, m_1, m_5, m_2] \\
L_{m_5} = [f_1, f_3, f_4, f_5, f_6] & L_{f_5} = [m_4, m_2, m_1, m_6, m_5, m_3] \\
L_{m_6} = [f_3, f_5, f_6, f_1, f_4] & L_{f_6} = [m_3, m_6, m_2, m_5, m_1, m_4] \\
Q_M = [2, 3, 1, 2, 1, 3]^T & Q_F = [2, 2, 2, 2, 2, 2]^T.
\end{array}$$

A male–female pair, (m, f) is considered a *feasible* pair if m and f are in each other's preference lists. That is, $m \in L_f$ and $f \in L_m$. A *matching* is defined to be a set of feasible male–female pairs $\{(m, f)\}$ such that $\forall m \in M$, m appears in at most q_m pairs and $\forall f \in F$, f appears in at most q_f pairs. A matching γ is said to be *stable* if any feasible pair $(m, f) \notin \gamma$ implies that at least one of the two persons, m or f is matched to its full quota of partners all of whom he or she considers better. This implies that for any stable matching γ, there cannot be any unmatched feasible pair (m, f) that can be paired with both m and f becoming better off. Let γ_m denote the set of partners of male m under the stable matching γ. The partners of female f are denoted by γ_f.

Let Γ denote the set of all stable matchings γ for a given instance of the multiple partner stable matching problem $P = (M, F, L_M, L_F, Q_M, Q_F)$. Accordingly, Γ_m is the set of all possible sets of partners γ_m that the male m can have in different stable matchings in Γ. For a female f, Γ_f is similarly defined. Given an instance of the many-to-many stable matching problem, γ_m^M defines the set of females assigned to male m in the male-optimal stable matching.

3. Notion of optimality

The optimality criterion for one-to-one matching (by McVitie and Wilson [13] and Irving et al. [9]) stipulates that the sum of ranks of partners for all males and females be minimized, as every person unambiguously prefers a partner of lower rank over a partner of higher rank. In a multiple partner context, things get complicated because a person's rankings of individuals alone may not be sufficient to determine his or her preference orderings over combinations of them. For example, (i) a person who wants up to two matches need not be indifferent to the partner combinations (1, 6), (2, 5) and (3, 4) (where the numbers denote the partners' ranks in the person's preference list), (ii) he or she may actually prefer (1, 6) over (2, 4) due to an overbearing preference for the partner ranked 1 even though (1, 6) has a greater sum of ranks than (2, 4), and (iii) the person may prefer (2, 6) over (2, 5) if partners ranked 2 and 6 are complementary and their combination is of greater value to him/her than the combination (2, 5).

Out of the three complications presented in the above illustration, the first two can be avoided using weighted preference lists. However, such an approach would come at a significant cost because (i) an individual needs to compute mutually consistent weights that capture his or her preference ordering over all acceptable partner-combinations, and (ii) the weights need to be normalized across individuals because the optimality criterion would consider the sum of weights across all individuals in the matching.

We show that the many-to-many stable matching has the necessary structure to explicitly rule out the above complications if a *no-complementarities* condition is imposed on preference orderings of males and females over combinations of partners. The no-complementarities condition (for males) states: given two sets of partners, A_1 and A_2 ($A_1, A_2 \subset F$), if a male m prefers A_1 at least as much as A_2, and m strictly prefers f_1 over f_2 ($f_1, f_2 \in F \setminus (A_1 \cup A_2)$), then m strictly prefers $A_1 \cup \{f_1\}$ over $A_2 \cup \{f_2\}$. This assumption is similar to the substitutability assumption widely used in the literature (For example: Roth [17] and Martinez et al. [12]) and can be derived from it. The assumption is intuitive and is plausible in most situations except those where specific partners may be strong complements of one-another.

Given a set of partners γ_m in a stable matching γ, we define the *dissatisfaction score* $DS(\gamma_m)$ of male m to be the sum of position numbers (or ranks) in its preference list L_m for the members in γ_m. $DS(\gamma_m) = \sum_{f \in \gamma_m} R_m(f)$ where $R_m(f)$ is the rank given by male m to the female f. The dissatisfaction score $DS(\gamma_f)$ for a female f is similarly defined. The dissatisfaction score of a matching γ is defined as the sum of the dissatisfaction scores of all persons involved. We show that the dissatisfaction score as defined above and the no-complementarities condition stated earlier impose a strict ordering on a person's preferences over all possible set of partners that he or she may have in any stable marriage for any given instance of the problem $P = (M, F, L_M, L_F, Q_M, Q_F)$.

First, we note some useful properties of the many-to-many stable marriage problem (due to Baiou and Balinski [3]). The results below are stated for males. They are also true for females due to the symmetrical nature of the problem.

Property 1. *A male $m \in M$ is assigned the same number of partners, $N_P(m)$, in all stable matchings. Further, if $N_P(m) < q_m$, then m has the same set of partners in all stable matchings.*

Property 2. *Suppose γ and γ^* are stable matchings that assign different sets of partners to a male m. Then there is one (say γ) such that if $(m, f) \in \gamma$ and $(m, f^*) \in \gamma^* \setminus \gamma$, then $R_m(f) < R_m(f^*)$.*

A useful corollary of Property 2 is that if a person is assigned different sets of partners in different stable matchings, then his or her least preferred partner in each of them must be different. Let the function min specify the least preferred partner of a person amongst his or her given set of partners. Accordingly, $\min(\gamma_m)$ and $\min(\gamma_f)$ specify the least preferred partners of male m and female f in the stable matching γ.

The Properties 1 and 2 stated above and the no-complementarities condition lead us to the following important result (stated for males):

Theorem 1. *Suppose γ_m and γ_m^* are two distinct sets of partners of the male m under the stable matchings γ and γ^* respectively. Then,* (i) *$DS(\gamma_m) \neq DS(\gamma_m^*)$, and* (ii) *If $DS(\gamma_m) < DS(\gamma_m^*)$ then m prefers γ_m over γ_m^* and vice versa.*

Proof. By Property 1, m must be matched to the same number of females in all stable matchings. Without loss of generality, Property 2 allows us to assume that $R_m(\min(\gamma_m)) < R_m(\min(\gamma_m^*))$ (implying that $\min(\gamma_m^*)$ is to the right of $\min(\gamma_m)$ in male m's preference list L_m). Each female in L_m to the right of $\min(\gamma_m)$ corresponds to at most one set of partners for m (amongst which it is the least preferred partner).

We consider the first female to the right of $\min(\gamma_m)$ in L_m which is the least preferred partner of m in some stable marriage, say $\gamma^{**} \in \Gamma$. By Property 2, we note that any female $f \in \gamma_m^{**} \setminus \gamma_m$ must be to the right of any female $f' \in \gamma_m$ in L_m. Further, since $|\gamma_m| = |\gamma_m^{**}|$, each such female f is a replacement of some other female $f' \in \gamma_m$ which was to the left of f in the preference list L_m. Each replacement of f' with f leads to an increase in the dissatisfaction score of m so that $DS(\gamma_m) < DS(\gamma_m^{**})$. Further, the no-complementarities assumption applied successively to each such replacement implies that m strictly prefers γ_m over γ_m^{**}.

If γ_m^{**} is identical to γ_m^*, it completes the proof. Else continue the above step until $\gamma_m^{**} = \gamma_m^*$. □

Theorem 1 is significant because it shows that we can obtain a strict preference ordering over all possible sets of acceptable stable marriage partners for any person by specifying only the preference orderings on individuals and imposing a no-complementarities assumption on the preferences over combinations. This obviates the need for preference functions which *apriori* specify the orderings over all possible subsets of members of the opposite sex which would be of exponential size. We can now use the terms *better* or *worse* unambiguously to compare any two sets of stable marriage partners of a person and we can do so by comparing either the least preferred partners or the dissatisfaction scores.

Theorem 1 allows us to propose that the minimization of the sum of dissatisfaction scores across all persons can be used as an egalitarian measure of optimality for the many-to-many stable marriage problem which is specified by $P = (M, F, L_M, L_F, Q_M, Q_F)$ for which the preferences over combinations of partners additionally satisfy the no-complementarities condition. Formally,

$$\textit{Minimize}_{\gamma \in \Gamma} \quad : \sum_{m \in M} DS(\gamma_m) + \sum_{f \in F} DS(\gamma_f). \tag{1}$$

The optimality criterion can also be restated in terms of matched pairs using the definition of the dissatisfaction score.

$$Minimize_{\gamma\in\Gamma} \quad : \sum_{(m,f)\in\gamma} [R_m(f) + R_f(m)]. \tag{2}$$

We note that the above optimality measure is also the natural generalization of the one proposed for the one-to-one marriage problem by McVitie and Wilson [13] for which a polynomial time algorithm was provided by Irving et al. [9].

The treatment of optimality using dissatisfaction score can be generalized to include weighted ranks which allows persons to specify their preferences more accurately, though with the disadvantages discussed earlier. The results presented in this paper are also true for the weighted rank preferences. The only requirement is that the preference orderings on individuals be strict. We will henceforth refer only to the case with unity weights for simplicity of exposition.

4. Reduction of search space

Irving and Leather [8] and Irving et al. [9] describe a methodology for the single partner stable marriage problem in which starting with the male-optimal solution, all the stable matchings can be generated by successive elimination of what they call *rotations*. They conclude that the enumeration problem is #P-complete and provide a polynomial time algorithm for finding a matching which satisfies the egalitarian optimality criterion.

The basis of their methodology lies in the process by which the *rotations* get exposed and eliminated. A rotation is a cycle comprising r male–female pairs (m_i, f_i) such that f_i is m_i's current match and f_{i+1} ($i+1$ is taken modulo r) is the second in m_i's current list. Since every individual is matched to only one partner, it is convenient to eliminate all females to the left to the current match of a male and all males to the right of the current match of a female from the male-optimal solution. (By definition of the male-optimal solution, these cannot occur in any stable matching.) This step is sufficient to ensure that at least one rotation gets exposed as long as the female-optimal solution is not reached. The exposed rotation is then eliminated (m_i gets paired with f_{i+1}) and the process continued till all rotations are eliminated.

Let us examine what happens when we try to directly apply the above methodology to the many-to-many context — it is known (due to Baiou and Balinski [3]) that the corresponding male-optimal stable matching γ^M always exists and can be found in $O(n^2)$ steps. Similar to its one-to-one counterpart, γ^M has the property that there is no other stable matching in which any of the males is better off (has a lower dissatisfaction score) or any of the females worse off (has a greater dissatisfaction score).

Example 2. For the problem instance in Example 1, the male-optimal solution is: $\gamma^M = \{(m_1, f_2), (m_1, f_6), (m_2, f_1), (m_2, f_2), (m_2, f_4), (m_3, f_5), (m_4, f_4), (m_4, f_3), (m_5, f_1), (m_6, f_3), (m_6, f_5), (m_6, f_6)\}$.

Beyond the male-optimal solution, applying the methodology to the many-to-many case is far from straightforward. Consider γ^M: a male m may be matched to multiple and non-contiguous females in its preference list L_m. The status of the females (not matched to m) in L_m who fall between the matched ones in L_m is not immediately clear. Indeed, we show later that all such females can be deleted from the list L_m from further consideration; on the other hand the males not matched to a female f but who are lying between the matched ones in L_f need to be retained as they can occur in stable marriages yet to be identified. This illustrates why the pruning (or elimination) step is tricky — not removing the necessary entries may result in no rotations being exposed while unnecessary deletion may lead to some stable marriages not being found.

Another problem is in defining a meaningful rotation. In the single partner case, a male getting matched to its second-most-preferred partner constitutes logical *atomic* progress from the male-optimal towards the female-optimal solution. In the multiple partner setup, such a progress may be possible in multiple ways — a subset of a male's current partners may be swapped with another subset such that the dissatisfaction score of the male increases by one. It is easy to see that the number of such choices is exponential.

A key contribution of this paper is to define a generalized concept of rotation and show how they can be exposed and eliminated. This is achieved through the definitions that follow.

Definition 1 (*Initial Pruning*). Given the male-optimal solution γ^M of an instance P, the initial pruning step consists of: (a) removing the females $f \notin \gamma_m^M$ from each male list L_m for which $R_m(f) < R_m(\min(\gamma_m^M))$, (b) removing the males $m \notin \gamma_f^M$ from a female list L_f for which $R_f(m) > R_f(\min(\gamma_f^M))$, (c) removing f from m's list if the (so far) reduced list of f does not contain m, and (d) removing m from f's list if the (so far) reduced list of m does not contain f.

Next, we introduce the concept of a *meta-rotation*, which is a many-to-many generalization of the concept of a *rotation*. The definition of a meta-rotation is motivated by the observation that if a person becomes better off or worse off, his or her least preferred partner (or min) must change.

Definition 2 (*Meta-Rotation*). Given a problem instance P, a meta-rotation ρ is defined as an ordered sequence of feasible male–female pairs $\{(m_0, f_0), \ldots, (m_i, f_i), \ldots, (m_{r-1}, f_{r-1})\}$, $r \geq 2$ such that $f_{(i+1)\ modulo\ r} = \text{smin}(\gamma_{m_i}^M)$ and $m_i = \min(\gamma_{f_i}^M)$. Here, $\text{smin}(\gamma_{m_i}^M)$ denotes the female to the immediate right of $\min(\gamma_{m_i}^M)$ in L_{m_i}. Such a meta-rotation is said to be exposed in P relative to its male-optimal stable matching, γ^M.

Definition 3 (*Meta-Rotation Elimination*). A meta-rotation ρ exposed in P, is said to be eliminated when for each $(m_i, f_i) \in \rho, m_i$ gets matched to $f_{(i+1)\ modulo\ r}$ in place of f_i and for each female f_i, its preference list L_{f_i} is modified to delete all males to the right of its new least preferred partner, and the preference lists of the males are correspondingly modified to remove females that are no longer their partners in feasible pairs.

Definition 4 (*R-instance*). A problem instance $\hat{P} = (M, F, \hat{L}_M, \hat{L}_F, Q_M, Q_F)$ is defined to be a reduced instance (or R-instance) of the problem instance $P = (M, F, L_M, L_F, Q_M, Q_F)$ if it is obtained either by applying Initial Pruning on P or by the Elimination of a Meta-rotation from another R-instance $\hat{P}^*$ of P.

Example 3. Application of initial pruning to the problem instance P in Example 1 yields the R-instance $\hat{P}$ with the following male and female lists (the matched pairs corresponding to the male-optimal solution in Example 2 are in boldface):

$$\begin{array}{llll}
\hat{L}_{m_1} = [\mathbf{f_2}, \mathbf{f_6}, f_1, f_3, f_5, f_4] & \hat{L}_{f_1} = [\mathbf{m_5}, m_1, \mathbf{m_2}] \\
\hat{L}_{m_2} = [\mathbf{f_1}, \mathbf{f_2}, \mathbf{f_4}, f_5, f_6, f_3] & \hat{L}_{f_2} = [\mathbf{m_1}, \mathbf{m_2}] \\
\hat{L}_{m_3} = [\mathbf{f_5}, f_3, f_4, f_6] & \hat{L}_{f_3} = [m_3, \mathbf{m_4}, m_5, m_1, m_2, \mathbf{m_6}] \\
\hat{L}_{m_4} = [\mathbf{f_4}, \mathbf{f_3}, f_5] & \hat{L}_{f_4} = [m_6, m_3, \mathbf{m_4}, m_1, m_5, \mathbf{m_2},] \\
\hat{L}_{m_5} = [\mathbf{f_1}, f_3, f_4, f_5, f_6] & \hat{L}_{f_5} = [m_4, m_2, m_1, \mathbf{m_6}, m_5, \mathbf{m_3}] \\
\hat{L}_{m_6} = [\mathbf{f_3}, \mathbf{f_5}, \mathbf{f_6}, f_4] & \hat{L}_{f_6} = [m_3, \mathbf{m_6}, m_2, m_5, \mathbf{m_1}].
\end{array}$$

The meta-rotation $\rho_1 = \{(m_3, f_5), (m_6, f_3), (m_2, f_4)\}$ is exposed in $\hat{P}$.

Note that every R-instance $\hat{P}$ of P is a stable marriage problem instance in its own right and therefore has meta-rotations defined on it unless the male optimal matching coincides with the female optimal matching. Since the preference lists in $\hat{P}$ are a subset of those in P, it is clear from the definition of stability that a stable matching in P, if it can be defined in $\hat{P}$, will also be stable in $\hat{P}$.

With the help of the foregoing definitions, we can now show that if we start with the male-optimal stable matching γ^M of P and successively identify and eliminate the meta-rotations, we would finally reach the female-optimal stable matching γ^F of P. If the process is carried out exhaustively for all possible sequences of meta-rotation eliminations, it would generate all possible stable marriages for the problem instance P.

We now obtain the intermediate results required for the purpose. We start by showing that the male–female pairs eliminated by initial pruning do not occur in any stable matching.

Lemma 1. *The sets of all possible stable matchings for the R-instance obtained by initial pruning of P is the same as Γ, set of stable matchings for P.*

Proof. For step (a) of initial pruning, we note by Property 2 that if $(m, f) \in \gamma \setminus \gamma^M$, then $R_m(f) > R_m(\min(\gamma_m^M))$. Therefore, a female f which is not paired with m in γ^M and is to the left of $\min(\gamma_m^M)$ cannot be paired with m in any stable matching. For males deleted from a female f's list in (b), we note that they are preferred less by f than her least preferred partner in her worst possible set of partners and hence cannot be paired with f in any stable matching.

Steps (c) and (d) only remove infeasible pairs. We have thus shown that pairs removed from consideration by initial pruning cannot occur in any stable matching for the problem instance P. Therefore, all stable matchings of P can also be found in $\hat{P}$. Further, removal of some pairs cannot introduce any new stable matching. Hence the result. □

Next, we show that every R-instance of P has a corresponding male-optimal stable matching which is also a stable matching for original instance P.

Lemma 2. *An R-instance $\hat{P}$ of the stable marriage problem P has a male-optimal stable matching $\hat{\gamma}^M$ in which each male m is matched to the first $N_P(m)$ females in his list $\hat{L}_m$ and the least preferred partner for each female f is the right-most in her list $\hat{L}_f$. Further, $\hat{\gamma}^M \in \Gamma$, the set of stable matchings for P.*

Proof. The proof is by induction. We note that an R-instance $\hat{P}$ can be generated from P by initial pruning and then by successive meta-rotation eliminations. For the R-instance of P obtained by initial pruning, Lemma 1 implies that the male-optimal stable matching γ^M for P is indeed the required male-optimal stable matching of the R-instance.

Let $\hat{P}^*$ be an R-instance of P which satisfies Lemma 2. Therefore, $\hat{\gamma}^{*M} \in \Gamma$. Consider an R-instance $\hat{P}^{**}$ obtained by eliminating an exposed meta-rotation ρ from $\hat{P}^*$. Clearly each female f's least preferred partner occurs last in her list $\hat{L}^{**}_f$ by the definition of meta-rotation elimination. A male m who is not in the meta-rotation ρ continues to be matched to the first $N_P(m)$ females in his list $\hat{L}^{**}_m$. A male m who is a part of ρ ends up pairing with the female immediately to the right of his least preferred match in lieu of a female whom he preferred more. The latter is removed from his list. Thus, m continues to be matched to the first $N_P(m)$ females in his new list $\hat{L}^{**}_m$.

Note that the number of partners of each male and female does not change in a meta-rotation elimination. Therefore, the newly updated set of matched pairs obtained by eliminating ρ continues to be a matching for P. We next show that it constitutes a stable matching for the R-instance $\hat{P}$, obtained by initial pruning of P.

For contradiction, suppose that it does not constitute a stable matching for $\hat{P}$. Then there exists a pair (m, f) not in this matching where $m \in \hat{L}_f$ and $f \in \hat{L}_m$ and m and f prefer each other over their respective least preferred partners in this matching. Consider the male m. Since his current partners occupy the first $N_P(m)$ positions in his current list $\hat{L}^{**}_m$, therefore, f must have been deleted from his old list $\hat{L}_m$ in a meta-rotation elimination step. At that time, f considered m worse than its least preferred partner. Since meta-rotation eliminations can only improve the least preferred partners of females, f must consider m worse than its current least preferred partner. This contradicts our assumption. Therefore, the matching (say $\hat{\gamma}^{**}$) is a stable matching for $\hat{P}$. By Lemma 1, it is also a stable matching for P. Further, since $\hat{\gamma}^{**}$ is contained within the preference lists of $\hat{P}^{**}$, it is also a stable matching for $\hat{P}^{**}$. Note that it is also the required male optimal stable matching $\hat{\gamma}^M$ for $\hat{P}^{**}$ since each male m is matched to its first $N_P(m)$ partners in its list $\hat{L}^{**}_m$. □

Example 4. Lemma 2 is easily verified in Example 3 by observation. The male-optimal stable matching (shown in boldface) for the R-instance obtained by initial pruning corresponds to the male optimal stable matching for P.

Further elimination of meta-rotation $\rho_1 = \{(m_3, f_5), (m_6, f_3), (m_2, f_4)\}$ from the R-instance in Example 3 gives an R-instance for which the male optimal solution $\hat{\gamma}^M$ is: $\{(m_1, f_2), (m_1, f_6), (m_2, f_1), (m_2, f_2), (m_2, f_5), (m_3, f_3), (m_4, f_4), (m_4, f_3), (m_5, f_1), (m_6, f_5), (m_6, f_6), (m_6, f_4)\}$. It can be checked that this also constitutes a stable matching for P.

We will now show that if there is a stable matching for P in which a male is worse off than in the male-optimal stable matching corresponding to an R-instance of P, then there must exist a meta-rotation exposed in that R-instance.

Lemma 3. *If m and f are deleted from each other's preference lists on the elimination of meta-rotation ρ exposed in $\hat{P}$ relative to its male-optimal stable matching $\hat{\gamma}^M$, and $(m, f) \notin \hat{\gamma}^M$, then (m, f) cannot occur in any stable matching γ of P.*

Proof. Suppose $(m, f) \notin \hat{\gamma}^M$. Then let m_1 be the least preferred partner of f after eliminating ρ from $\hat{P}$. Then, f prefers m_1 to m. Let f_1 be the least preferred partner of m in $\hat{\gamma}^M$. Clearly, $f \neq f_1$. If $R_m(f) < R_m(f_1)$ then by Property 2, (m, f) cannot belong to any stable marriage of P. On the other hand, if $R_m(f) > R_m(f_1)$, then if (m, f) is a pair in some stable matching γ of P, then both m and f are better off in the stable matching $\hat{\gamma}^M$ than they are in γ which cannot be true. □

At this stage, it is useful to state an important property of the multiple partner stable marriage problem (due to Baiou and Balinski [3]).

Property 3. *Suppose γ and γ^* are stable matchings that assign different sets of partners to a male m. Then, $\{R_m(\min(\gamma_m)) < R_m(\min(\gamma_m^*))\} \Rightarrow \{R_f(\min(\gamma_f)) > R_f(\min(\gamma_f^*))\}$ for all f such that $(m, f) \in (\gamma \setminus \gamma^*) \bigcup (\gamma^* \setminus \gamma)$.*

Property 3 leads to the conclusion that if any pair $(m, f) \in \gamma$, for some stable matching $\gamma \in \Gamma$, it is not possible for both m and f to be simultaneously worse off (or simultaneously better off) in any stable matching γ' $((m, f) \notin \gamma')$ than they are in γ.

Lemma 4. *Given an R-instance $\hat{P}$ of the stable marriage problem P, with its male-optimal stable matching $\hat{\gamma}^M$, if there is a male m for whom $R_m(\min(\hat{\gamma}_m^M)) < R_m(\min(\gamma_m))$ for some stable marriage $\gamma \in \Gamma$ of P, then there is at least one meta-rotation exposed in $\hat{P}$.*

Proof. Suppose that there is a male m for whom $R_m(\min(\hat{\gamma}_m^M)) < R_m(\min(\gamma_m))$ for some stable marriage $\gamma \in \Gamma$ of P. We construct a sequence $\{(m_i, f_i)\}$ as follows: Let $m_0 = m$. By Lemma 1 and Lemma 3, no females to the right of $\min(\hat{\gamma}_{m_0}^M)$ which can be paired with m_0 in some stable matching get deleted during initial pruning and meta-rotation elimination, therefore, m_0 is worse off in γ implies that $\text{smin}(\hat{\gamma}_{m_0}^M)$ is defined in $\hat{L}_{m_0}$. Let $f_1 = \text{smin}(\hat{\gamma}_{m_0}^M)$. In γ, m_0 either gets matched to f_1 or to someone further right in his list. In the former case, f_1 becomes better off in γ because m_0 is to the left of her least preferred partner in $\hat{\gamma}^M$. In the latter case, f_1 must be matched to its full quota of partners, each of whom she prefers over m_0 otherwise the matching γ would not be stable. In this case too, f_1 becomes better off in γ than in $\hat{\gamma}^M$.

Now, f_1 is better off in γ implies that $R_{f_1}(\min(\gamma_{f_1})) < R_{f_1}(\min(\hat{\gamma}_{f_1}^M))$. Let $m_1 = \min(\hat{\gamma}_{f_1}^M)$. Clearly, $(m_1, f_1) \notin \gamma$ as all of f_1's partners in γ are preferred by her over m_1. Since f_1 is better off in γ, and (m_1, f_1) are partners in $\hat{\gamma}^M$ which by Lemma 2 is a stable matching of P, therefore, m_1 must be worse off in γ using Property 3.

We continue to build the chain where $f_{i+1} = \text{smin}(\hat{\gamma}_{m_i}^M)$ and $m_i = \min(\hat{\gamma}_{f_i}^M)$. We cannot progress indefinitely, so the sequence $\{(m_i, f_i)\}$ must cycle. Thus, we have constructively shown the existence of a meta-rotation in $\hat{P}$ (relative to the matching $\hat{\gamma}^M$). □

We can designate this meta-rotation as the meta-rotation generated in $\hat{P}$ by a male m which gets a worse set of partners in the new stable matching after eliminating the meta-rotation.

Lemma 5. *If $\{(m_0, f_0), \ldots, (m_{r-1}, f_{r-1})\}$, $r \geq 2$ is a meta-rotation exposed in some R-instance $\hat{P}$ of P relative to its male-optimal matching $\hat{\gamma}^M$ and in some stable matching $\gamma \in \Gamma$, $R_{m_k}(\min(\gamma_{m_k})) > R_{m_k}(\min(\hat{\gamma}_{m_k}^M))$ for a particular male m_k, then for each male m_i, $i \in 0 \ldots r-1$ in the meta-rotation, $R_{m_i}(\min(\gamma_{m_i})) > R_{m_i}(\min(\hat{\gamma}_{m_i}^M))$.*

Proof. On the lines of the proof for Lemma 4, we note that if m_k becomes worse off in γ than in $\hat{\gamma}$, then the female $f_{(k+1) \text{ modulo } r}$ should become better off. Continuing further, this should make the male $m_{(k+1) \text{ modulo } r}$ worse off. The Lemma 5 therefore follows. □

Corollary 1. *If $\{(m_0, f_0), \ldots, (m_{r-1}, f_{r-1})\}$, $r \geq 2$ is a meta-rotation exposed in some R-instance $\hat{P}$ of P relative to its male-optimal matching $\hat{\gamma}^M$ and in some stable matching $\gamma \in \Gamma$, $R_{f_k}(\min(\gamma_{f_k})) < R_{f_k}(\min(\hat{\gamma}_{f_k}^M))$ for a particular female f_k, then for each female f_i, $i \in 0 \ldots r-1$ in the meta-rotation, $R_{f_i}(\min(\gamma_{f_i})) < R_{f_i}(\min(\hat{\gamma}_{f_i}^M))$.* □

We are now in a position to show that every stable marriage for the problem instance P can be obtained as the male-optimal solution for some R-instance $\hat{P}$ of P.

Theorem 2. *Given a stable matching $\gamma_0 \in \Gamma$ for P, γ_0 is identical to the male-optimal stable matching $\hat{\gamma}^M$ for some R-instance $\hat{P}$ of P.*

Proof. Consider the R-instance $\hat{P}$ obtained by applying initial pruning to P. If $\gamma_0 = \gamma^M$ then by Lemma 1, $\gamma_0 = \hat{\gamma}^M$ also. Therefore, $\hat{P}$ is the required R-instance.

Suppose $\gamma_0 \neq \hat{\gamma}^M$. This implies that there is at least one male m which is worse off in γ_0 than in $\hat{\gamma}^M$. By Lemma 4, there exists a meta-rotation ρ exposed in $\hat{P}$. We also note (by proof methodology of Lemma 4) that ρ can be generated by starting from m so that m is included in the meta-rotation. We can eliminate ρ from $\hat{P}$ to yield a new R-instance $\hat{P}^*$. By Lemma 5, we note that in any stable matching $\gamma \in \Gamma$ (including γ_0), in which $R_m(\min(\gamma_m)) > R_m(\min(\hat{\gamma_m}^M)$, the least preferred partner of every male included in ρ must also be worse than in $\hat{\gamma}^M$. By Lemma 3, a meta-rotation elimination does not remove any female right of the least preferred partner of a male from his list unless they cannot be paired in any stable marriage, therefore, it is ensured that for any male $m_i \in M$, $R_{m_i}(\min(\hat{\gamma}_{m_i}^{M*})) \leq R_{m_i}(\min(\gamma_{0\ m_i}))$ which states that m_i prefers its least preferred partner in γ^{M*} at least as much as the least preferred partner in γ_0.

Suppose now, $R_{m_i}(\min(\hat{\gamma}_{m_i}^{M*})) = R_{m_i}(\min(\gamma_{0\ m_i}))$ for all $m_i \in M$. By Lemma 2, we know that $\hat{\gamma}^{M*} \in \Gamma$, the set of stable matchings for P. Since the least preferred partners uniquely define a person's set of partners, it follows that $\hat{\gamma}^{M*} = \gamma_0$ and $\hat{P}^*$ is the required R-instance.

Otherwise, there is at least one male m' which is worse off in γ_0 than in $\hat{\gamma}^{M*}$. We can again apply Lemma 4 to find a meta-rotation exposed in $\hat{P}^*$. This process terminated only when we get a R-instance such that for any $m_i \in M$, $R_{m_i}(\min(\hat{\gamma}_{m_i}^{M*})) = R_{m_i}(\min(\gamma_{0\ m_i}))$. At that point, the male-optimal stable matching for the R-instance is identical to γ_0. □

Therefore, every stable matching $\gamma \in \Gamma$ for P can be obtained by successive application of meta-rotation eliminations on the R-instance obtained by initial pruning of P. Since the male-optimal stable matching $\hat{\gamma}^M$ is unique for a given R-instance $\hat{P}$, we have also established a one-to-one correspondence between the stable matchings γ for P and the R-instances $\hat{P}$ of P.

Example 5. In Example 3, the elimination of meta-rotation $\rho_1 = \{(m_3, f_5), (m_6, f_3), (m_2, f_4)\}$ leads to the meta-rotation $\rho_2 = \{(m_1, f_6), (m_2, f_1)\}$ becoming exposed. Elimination of ρ_2 gives the female-optimal solution $\gamma^F = \{(m_1, f_2), (m_1, f_1), (m_2, f_2), (m_2, f_5), (m_2, f_6), (m_3, f_3), (m_4, f_4), (m_4, f_3), (m_5, f_1), (m_6, f_5), (m_6, f_6), (m_6, f_4)\}$.

Define Ω to be the set of meta-rotations for the problem instance P. $\rho \in \Omega$ if and only if ρ is a meta-rotation exposed in some R-instance $\hat{P}$ of the problem instance P. We will now define a partial order on the set Ω. For that, we need one additional result which is stated below.

Lemma 6. *No pair (m, f), $m \in M$, $f \in F$ can belong to two different meta-rotations $\rho_1, \rho_2 \in \Omega$.*

Proof. Suppose that the pair (m, f) belongs to meta-rotations ρ_1 and ρ_2. Since $\rho_1 \neq \rho_2$, there is a pair (m', f') that belongs to ρ_1 and not ρ_2.

Let $\hat{P}$ be the R-instance in which ρ_2 is exposed and let $\hat{\gamma}^M$ be the corresponding male-optimal stable matching in $\hat{P}$. If $m_1 = \min(\hat{\gamma}_{f'}^M)$ it implies $R_{f'}(m') \leq R_{f'}(m_1)$. For if $R_{f'}(m_1) < R_{f'}(m')$ then elimination of ρ_1 would force $R_f(\min(\hat{\gamma}_f^M)) < R_f(m)$ by Corollary 1.

Let $\hat{P}^*$ be the R-instance and $\hat{\gamma}^{M*}$ be the corresponding stable matching obtained after eliminating ρ_2.

If $(m_1, f') \in \rho_2$ and given that $(m', f') \notin \rho_2$, then $m_1 \neq m'$, and therefore, $R_{f'}(m') < R_{f'}(m_1)$. Since $(m, f) \in \rho_2$ and $\hat{P}^*$ arises from the elimination of ρ_2, f will be better-off in $\hat{P}^*$ than in $\hat{\gamma}^M$ i.e. $R_f(\min(\hat{\gamma}_f^{M*})) < R_f(m)$. Lemma 3 points out the presence of m' in f''s list. Furthermore this presence is not affected by the elimination of ρ_2. Hence, $R_{f'}(m') \leq R_{f'}(\min(\hat{\gamma}_{f'}^{M*}))$. Therefore, the pairs (m, f), (m', f') of ρ_1 and the R-instance $\hat{P}^*$ contradict Corollary 1 because f gets better off but f' does not.

In the other case, when $(m_1, f') \notin \rho_2$ then elimination of ρ_2 does not affect the presence of m_1 in f''s list. Hence, f gets worse-off i.e. $R_f(\min(\hat{\gamma}_f^{M*})) < R_f(m)$ and $R_{f'}(m_1) < R_{f'}(\min(\hat{\gamma}_{f'}^{M*}))$. Proceeding exactly as in the previous case we have a contradiction of Corollary 1. □

When a meta-rotation ρ' must be eliminated before a meta-rotation ρ is exposed in some R-instance $\hat{P}$ then we say that $\rho' < \rho$ or ρ' is a predecessor of ρ. Since all meta-rotations are generated by successive elimination of meta-rotations from the R-instance obtained by initial pruning of P, it can be seen that the reflexive transitive closure $\leq$ of the predecessor relation establishes a partial order on the set Ω. We call this as the meta-rotation poset $\Psi, \leq$. A closed subset A in $\Psi, \leq$ is a subset of $\Psi, \leq$ such that $\forall\ \rho \in A, \rho' < \rho \Rightarrow \rho' \in A$.

Example 6. In Example 5, $\Omega = \{\rho_1, \rho_2\}$. Meta-rotation ρ_1 is a predecessor of ρ_2. There are two closed subsets, $A_1 = \{\rho_1\}$ and $A_2 = \{\rho_1, \rho_2\}$.

We can now state the final result. This generalizes the corresponding result for one-to-one stable matchings by Irving and Leather [8].

Theorem 3. *There is a one-to-one correspondence between the stable matchings of a given problem instance P and the closed subsets of the meta-rotation poset $\Psi, \leq$ defined over its set of meta-rotations Ω.*

Proof. Given a closed subset A of $\Psi, \leq$, we can eliminate all the meta-rotations present in A since all the predecessors of those meta-rotations are also in A. This yields an R-instance $\hat{P}$. By Lemma 2, the male-optimal stable matching of $\hat{P}$ is a stable matching for P. Given two different closed subsets $A_1 \neq A_2$, we would get two different stable matchings. This follows from Lemma 6 whereby different meta-rotations eliminate different pairs and so any two R-instances, $\hat{P} \neq \hat{P}^*$ can only be obtained by eliminating two different sets of meta-rotations. Hence, each closed subset corresponds to a unique stable matching.

Given a stable matching $\gamma \in \Gamma$, by Theorem 2, $\gamma = \hat{\gamma}^M$ for some R-instance $\hat{P}$ of P. The set of meta-rotations that are eliminated to get any R-instance $\hat{P}$ must be closed because if $\rho_1 < \rho_2$, then ρ_1 must be eliminated before ρ_2 can become exposed. Hence, γ corresponds to a closed subset of $\Psi, \leq$. Also, by Lemma 2, a stable matching $\gamma^* \neq \gamma$ would correspond to a different R-instance and therefore by Lemma 6, a different closed set of meta-rotations. □

5. An efficient algorithm for 'optimal' stable matching

The problem of finding an 'optimal' stable matching now boils down to constructing the meta-rotation poset and finding its closed subset that minimizes the sum of the dis-satisfaction scores.

A directed acyclic graph corresponding to the meta-rotation poset $\Psi, \leq$ can be constructed using the predecessor relationship defined on the meta-rotations. The vertices of this graph correspond to meta-rotations and directed edges are added as described below.

Consider a meta-rotation ρ exposed in the R-instance $\hat{P}$. If $(m_i, f_i) \in \rho$ then the male m_i must be the last in the current list of the female f_i otherwise ρ cannot be exposed. Consider the first male m to the right of m_i in L_{f_i} such that (m, f_i) is a member of some other meta-rotation (say ρ^*). From the definition of meta-rotation elimination, it follows that eliminating ρ^* would make m_i the last in the list of f_i. We create a directed edge from ρ^* to ρ.

Similarly, for the male lists, consider the male optimal stable matching $\hat{\gamma}^M$ for an R-instance $\hat{P}$. Let $(m_i, f_i) \in \rho$ where ρ is a meta-rotation exposed in $\hat{P}$. Let $f_1 = \min(\hat{\gamma}^M_{m_i})$ and $f_2 = \text{smin}(\hat{\gamma}^M_{m_i})$. Eliminating all females f for whom $R_{m_i}(f_1) < R_{m_i}(f) < R_{m_i}(f_2)$ will make f_1 and f_2 contiguous in the list of m_i. Note that the pairs (m_i, f) with f as defined above would not occur in any stable matching and the meta-rotations that cause them to be eliminated from the list of m_i are predecessors of ρ. We create a directed edge from each such meta-rotation to the meta-rotation ρ.

It is easy to verify that the resulting directed acyclic graph has the property that if ρ_1 is predecessor of ρ_2, then there is a directed path from ρ_1 to ρ_2. To see this, consider the situation when ρ_1 is exposed (at this time, ρ_2 is not not exposed because ρ_1 is yet to be eliminated). Since ρ_2 is not exposed, our construction would imply at least one directed edge into ρ_2. Consider the set of all directed edges into ρ_2. It must be the case that either one of these edges is from ρ_1 or else ρ_1 must be predecessor of at least one of the meta-rotations that has a directed edge (path) to ρ_2 (else one can eliminate all these edges without eliminating ρ_1). Proceeding thus, we have a directed path from ρ_1 to ρ_2. The graph therefore preserves all the closed subsets of the meta-rotation poset $\Psi, \leq$ and thus contains all the stable matchings for the problem.

Given a meta-rotation $\rho = \{(m_0, f_0), \ldots, (m_{r-1}, f_{r-1})\},\ r \geq 2$ define its weight w_ρ in the following manner: $w_\rho = \sum_{i=0}^{r-1} [R_{m_i}(f_i) + R_{f_i}(m_i) - R_{m_i}(f_{i+1}) - R_{f_i}(m_{i-1})]$ where $(i+1)$ and $(i-1)$ are taken modulo r. The weight of a closed subset A of the meta-rotation poset $\Psi, \leq$ is defined to be the sum of weights of the meta-rotations in A.

Note that the weight of a meta-rotation equals the reduction in the dis-satisfaction score as a result of eliminating the meta-rotation. In other words, if the stable matching γ' is obtained by eliminating ρ from the stable matching γ then $DS(\gamma') = DS(\gamma) - w_\rho$.

Let the dissatisfaction score for the male-optimal stable matching γ^M be DS_0. The dissatisfaction score of the stable matching obtained by eliminating meta-rotations of a closed subset $A = \{\rho_1 \ldots \rho_k\}$ would therefore equal $DS_0 - \sum_{i \in 1 \ldots k} w_{\rho_i}$. The stable matching corresponding to the maximum weight closed subset would have the least dissatisfaction score amongst all stable matchings for P.

Putting everything together, a stable matching γ^* that minimizes the sum of dissatisfaction scores for the instance $P = (M, F, L_M, L_F, Q_M, Q_F)$ can be found in polynomial time by the following steps:

(1) Obtain the male-optimal stable matching γ^M.
(2) Apply initial pruning to P to get an R-instance $\hat{P}$.
(3) Find a meta-rotation ρ exposed in $\hat{P}$ (if one exists); eliminate ρ.
(4) Iterate previous step until no such ρ can be found.
(5) Construct the directed acyclic graph for the weighted meta-rotation poset $\Psi, \leq$.
(6) Identify maximum-weight closed subset A of $\Psi, \leq$.
(7) Eliminate the meta-rotations of A from the R-instance obtained by initial pruning of P. The male-optimal stable matching for this R-instance is the required optimal stable matching γ^*.

Example 7. In Example 5, $w(\rho_1) = 8$ and $w(\rho_2) = -2$ The maximum weight closed subset is therefore $A = \{\rho_1\}$. Therefore, the 'optimal' stable matching is: $\gamma = \{(m_1, f_2), (m_1, f_6), (m_2, f_1), (m_2, f_2), (m_2, f_5), (m_3, f_3), (m_4, f_4), (m_4, f_3), (m_5, f_1), (m_6, f_5), (m_6, f_6), (m_6, f_4)\}$. It is obtained by eliminating ρ_1 from the R-instance obtained by initial pruning of P.

We now determine the complexity of the algorithm. Let $n = max(|M|, |F|)$. Generating the male-optimal matching (Step 1) takes $O(n^2)$ time (Baiou and Balinski [3]). Since $\forall\, m \in M,\ q_m \leq |F|$ and $\forall\, f \in F,\ q_f \leq |M|$ the initial pruning (Step 2) requires at most $O(|M| * |F|^2 + |F| * |M|^2)$ steps which is bounded by $O(n^3)$. A meta-rotation, if exists, can be identified and eliminated in $O(n^2)$ steps in an iteration of Step 3. Note that as pointed out earlier, a pair (m, f) can occur in only one meta-rotation. Since a meta-rotation must have at least two pairs and there are at most $(|M| * |F|)$ pairs to be eliminated, the number of iterations of Step 3 is bounded by $O(n^2)$. Hence Steps 3 and 4 together take $O(n^4)$ steps to eliminate all possible meta-rotations.

The directed acyclic graph for the poset $\Psi, \leq$ has no more than $O(n^2)$ edges because there can be at most n^2 distinct male–female pairs and each pair contributes at most one directed edge to the graph (this edge is from the meta-rotation that eliminates the pair to another meta-rotation of which the eliminating meta-rotation is a predecessor). To construct the graph, one can use the meta-rotations to label the pairs (in the male and female lists) left after initial pruning. For each meta-rotation ρ, label the pairs that are members of ρ, those that are otherwise eliminated by ρ and the females that are min and smin for each male in the context of ρ. Since the elimination of pairs by meta-rotations is contiguous in the female lists, the labels can be added in time proportional to number of pairs and therefore in $O(n^2)$ time over all rotations. The graph can then be constructed by a simple scan of the male and female lists. Step 5 therefore takes $O(n^2)$ time.

Once the poset has been constructed and weights assigned, its maximum weight closed subset can be found (Step 6) in $O(n^6)$ time. (The computation of maximum weight closed subset of a weighted poset (or directed acyclic graph) is a well known problem. See Picard [14], Rhys [16] and Picard and Queyranne [15] for $O(n^6)$ solutions using network flow and related algorithms.) The overall algorithm complexity is thus bounded by $O(n^6)$.

Before closing, we examine the complexity of counting the number of stable matchings for a given instance of the multiple partner stable marriage problem P. Any given matching can be checked for stability in polynomial time which implies that the enumeration problem is clearly in #P. The problem instance P where $q_m = 1,\ \forall\, m \in M$ and $q_f = 1,\ \forall\, f \in F$ is also an instance of the single partner stable marriage problem which has been shown to be #P-complete (Irving and Leather [8]). Therefore, determining the number of stable matchings for an instance of the multiple partner stable marriage problem is also #P-complete.

6. Concluding remarks

In this paper, we considered an egalitarian measure of optimality for the multiple partner stable marriage problem (with incomplete lists) and provided a sufficient condition under which the criterion could be applied. Further, we provided a polynomial time algorithm for obtaining a stable matching which satisfied the optimality criterion. By doing so, we generalized some of the results known for the corresponding one-to-one problem.

The polynomial complexity is significant because the problem of determining the number of all stable matchings for the problem is #P-complete. In the process of solving the problem at hand, we proposed a novel concept of *meta-rotations* which extends the earlier known concept of *rotations* (Irving and Leather [8]) and makes it useful as a search space reduction technique for search problems. We also showed that a useful property of the multiple partner stable marriage problem is that under a no-complementarities assumption on the preferences over combinations of partners, specifying the preference ordering over individuals alone is sufficient to ensure that the preferences of males and females turn out to be strictly ordered over all possible sets of partners that they can get in any stable matching. The results presented in this paper can accommodate weighted preferences of males and females so that the optimality criterion is not restrictive.

We note that the methodology to map stable matchings of the many-to-many problem to the antichains of a poset as well as the polynomial time algorithm to find the matching which minimizes the dissatisfaction score is generic and does not require the no-complementarities assumption. At the same time, the use of dissatisfaction scores as an egalitarian measure of optimality makes best sense when the assumption is used.

Acknowledgements

Our special thanks are due to Robert Irving for reviewing our work and providing encouraging feedback for taking it forward. We thank David Manlove, Robert Irving and Patrick Prosser for sharing some of their current work on the stable marriage problem. We are also thankful to Lloyd Shapley and K. Iwama for sharing their feedback on the work.

References

[1] A. Alkan, A class of multipartner matching markets with a strong lattice structure, Economic Theory 19 (4) (2002) 737–746.
[2] A. Alkan, On preferences over subsets and the lattice structure of stable matchings, Review of Economic Design 6 (2001) 99–111.
[3] M. Baiou, M. Balinski, Many-to-many matching: Stable polyandrous polygamy (or polygamous polyandry), Discrete Applied Mathematics 101 (2000) 1–12.
[4] D. Gale, L.S. Shapley, College admissions and the stability of marriage, American Mathematical Monthly 69 (1962) 9–15.
[5] D. Gale, M. Sotomayor, Some remarks on the stable matching problem, Discrete Applied Mathematics 11 (1985) 223–232.
[6] D. Gusfield, R.W. Irving, The Stable Marriage Problem: Structure and Algorithms, The MIT Press, Cambridge, 1989.
[7] G. Polya, R.E. Tarjan, D.R. Woods, Notes on Introductory Combinatorics, Birkhauser Verlag, Boston, Massachussets, 1983.
[8] R.W. Irving, P. Leather, The complexity of counting stable marriages, SIAM Journal of Computing 15 (3) (1986) 655–667.
[9] R.W. Irving, P. Leather, D. Gusfield, An efficient algorithm for the "optimal" stable marriage, Journal of the ACM 34 (3) (1987) 532–543.
[10] K. Iwama, D. Manlove, S. Miyazaki, Y. Morita, Stable marriage with incomplete lists and ties, Proceedings of ICALP (1999) 443–452.
[11] D.E. Knuth, Mariages Stables, Les Presses de l'Universite de Montreal, Montreal, 1976.
[12] R. Martinez, J. Masso, A. Neme, J. Oviedo, An Algorithm to Compute the Set of Many-to-many Stable Matchings, UAB.IAE Working papers 436, 2001.
[13] D. McVitie, L.B. Wilson, The stable marriage problem, Commucations of the ACM 114 (1971) 486–492.
[14] J. Picard, Maximum closure of a graph and applications to combinatorial problems, Management Science 22 (1976) 1268–1272.
[15] J. Picard, M. Queyranne, Selected applications of minimum cuts in networks, INFOR — Canadian Journal of Operations Research and Information Processing 20 (1982) 394–422.
[16] J. Rhys, A selection problem of shared fixed costs and network flows, Management Science 17 (1970) 200–207.
[17] A. Roth, Stability and polarization of interests in job matching, Econometrica 52 (1984) 47–57.
[18] A. Roth, M. Sotomayor, Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis, in: Econometrica Society Monographs, vol. 18, Cambridge University Press, 1990.
[19] M. Sotomayor, The lattice structure of the set of stable outcomes of the multiple partners assignment game, International Journal of Game Theory 28 (1999) 567–583.
[20] L.G. Valiant, The complexity of computing the permanent, Theoretical Computer Science 8 (1979) 189–201.
[21] L.G. Valiant, The complexity of enumeration and reliability problems, SIAM Journal of Computing 8 (1979) 410–421.