Pareto optimality
in the Roommates Problem

David J. Abraham and David F. Manlove

Pareto optimality
in the Roommates problem

David J. Abraham[1*] and David F. Manlove[2†]

[1] *Computer Science Department, Carnegie-Mellon University, 5000 Forbes Ave, Pittsburgh PA 15213-3890, USA. Email:* `dabraham@cs.cmu.edu`.

[2] *Department of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK. Email:* `davidm@dcs.gla.ac.uk`.

Abstract. We consider Pareto optimal matchings as a means of coping with instances of the Stable Roommates problem (SR) that do not admit a stable matching. Given an instance I of SR, we show that the problem of finding a maximum Pareto optimal matching is solvable in $O(\sqrt{n\alpha(m,n)}m\log^{3/2}n)$ time, where n is the number of agents and m is the total length of the preference lists in I. By contrast we prove that the problem of finding a minimum Pareto optimal matching is NP-hard, though approximable within 2. We also show that the problem of finding a Pareto optimal matching with the fewest number of blocking pairs is NP-hard. However, for a fixed integer K, we give a polynomial-time algorithm that constructs a Pareto optimal matching with at most K blocking pairs, or reports that no such matching exists.

1   Introduction

The Stable Roommates problem (SR) is a classical combinatorial problem that has been studied extensively in the literature [9, 13, 11, 20, 15, 14]. An instance I of SR contains a graph $G = (A, E)$ where $A = \{a_1, \ldots, a_n\}$ and $m = |E|$. We assume that G contains no isolated vertices. We interchangeably refer to the vertices of G as the *agents*, and we refer to G as the *underlying graph* of I. The vertices adjacent to a given agent $a_i \in A$ are the *acceptable* agents for a_i, denoted by A_i. If $a_j \in A_i$, we say that a_i *finds a_j acceptable*. Moreover we assume that in I, a_i has a linear order over A_i, which we refer to as a_i's *preference list*. If a_j precedes a_k in a_i's preference list, we say that a_i *prefers a_j to a_k*.

Let M be a matching in G and let $a_i \in A$. If $\{a_i, a_j\} \in M$ for some $a_j \in A$, we say that a_i is *matched* in M and $M(a_i)$ denotes a_j, otherwise a_i is *unmatched* in M. A *blocking pair* with respect to a matching M is an edge $\{a_i, a_j\} \in E \backslash M$ such that

---

(i) either a_i is unmatched in M, or a_i is matched in M and prefers a_j to $M(a_i)$, and (ii) either a_j is unmatched in M, or a_j is matched in M and prefers a_i to $M(a_j)$. A matching is *stable* if it admits no blocking pair.

Gale and Shapley [9] showed that an instance of SR need not admit a stable matching. Irving [13] gave an $O(m)$ algorithm that finds a stable matching or reports that none exists, given an instance I of SR. The algorithm in [13] assumes that in I, all preference lists are *complete* (i.e. $A_i = A\backslash\{a_i\}$ for each $a_i \in A$) and n is even, though it is straightforward to generalize the algorithm to the problem model defined here (i.e. the case of *incomplete lists*) [11]. Henceforth we denote by SRC the special case of SR in which all preference lists are complete.

Empirical results [19] seem to suggest that, as n increases, the probability that an SR instance with n agents admits a stable matching decreases fairly steeply. For example, for various values of n, Thomson [22] generated x_n random SRC instances, each with n agents, and calculated p_n, the proportion of instances that admitted a stable matching. For $n = 10$, 100, 1000 and 10000, the values of p_n were 90.1%, 65.3%, 37.7% and 18.7% respectively. In the first two cases $x_n = 10000$, whilst in the last two cases $x_n = 1000$.

Given this observation, it is natural to consider an alternative, weaker optimality property that could always be satisfied by some matching in an instance of SR. One such property is *Pareto optimality*. Informally, a matching M is Pareto optimal if there is no other matching M' such that some agent is better off in M' than in M, whilst no agent is worse off in M' than in M.

Formally, Pareto optimality may be defined as follows. Firstly we define the preferences of an agent over matchings. Given two matchings M and M', we say that an agent a_i *prefers M' to M* if either (i) a_i is matched in M' and unmatched in M, or (ii) a_i is matched in both M and M' and prefers $M'(a_i)$ to $M(a_i)$. Given this definition, we may define a relation $\prec$ on the set of all matchings as follows: $M' \prec M$ if and only if no agent prefers M to M', and some agent prefers M' to M. It is straightforward to show that $\prec$ forms a strict partial order over the set of matchings in I. A matching is defined to be *Pareto optimal* if and only if it is $\prec$-minimal. Intuitively a matching is Pareto optimal if no agent a_i can improve without another agent a_j becoming worse off.

Pareto optimality has received much attention, particularly from the Economics community, and has recently been considered in the bipartite graph setting from an algorithmic point of view [1]. As a further motivation for considering Pareto optimality, we note that, in many applications a prime objective is to match as many agents as possible. It is known that, given an instance I of SR, all stable matchings in I (assuming at least one exists) are of the same size [11, Theorem 4.5.2]. However in Section 2, we give an infinite family of SR instances for which Pareto optimal matchings may have different sizes, and moreover the size of a stable matching is half the size of a maximum cardinality Pareto optimal matching (henceforth a maximum Pareto optimal matching).

Given that Pareto optimal matchings may be of different sizes in I, it is natural to consider the algorithmic complexity of each of the problems of finding a maximum and minimum Pareto optimal matching. Moreover, if I does not admit a stable matching, it is of interest to ask whether there is an efficient algorithm for constructing a matching that contains the fewest number of blocking pairs, so as to

```
1 :  4  3  2              1 :   6  2  . . .
2 :  3  4  1              2 :   3  1  . . .
3 :  1  2  4              3 :   2  4  . . .
4 :  2  1  3              4 :   5  3  . . .
  Instance $I_1$          5 :   4  6  . . .
                          6 :   1  5  . . .
                            Instance $I_2$
```

Figure 1: Two instances of SRC.

obtain a matching that is "as stable as possible". As we demonstrate in Section 4, such a matching must be Pareto optimal.

The remainder of this paper is organised as follows. In Section 2 we give a necessary and sufficient condition for an arbitrary matching to be Pareto optimal, leading to an $O(m)$ algorithm for checking the Pareto optimality of a matching. We also give a straightforward $O(m)$ greedy algorithm for finding a Pareto optimal matching in an instance of SR. In Section 3 we show that the problem of finding a maximum Pareto optimal matching is solvable in $O(\sqrt{n\alpha(m,n)}m\log^{3/2}n)$ time, where α is the inverse of Ackermann's function. However in Section 4 we show that, given an instance of SR, the problem of finding a (Pareto optimal) matching with the fewest number of blocking pairs is NP-hard. On the other hand, for a fixed integer K, we give a polynomial-time algorithm that constructs a Pareto optimal matching with at most K blocking pairs, or else reports that no such matching exists. Finally, in Section 5, we consider minimum Pareto optimal matchings. We show that, given an SR instance, the problem of finding such a matching is NP-hard, though approximable within a factor of 2.

We conclude this section with two remarks regarding related work. Firstly, an alternative method has been considered in the literature for coping with instances of SR that do not admit a stable matching. Tan [21] presented an $O(n^2)$ algorithm that finds, given an SR instance I, a largest matching M in I with the property that the matched pairs in M are stable within themselves. However again we note that such a matching may be half the size of a maximum Pareto optimal matching. For example, we may choose any insoluble SRC instance I with 4 agents – there are 48 such instances [11, p.220]. In I, Tan's algorithm is bound to construct a matching of size 1, though the size of a maximum Pareto optimal matching is 2. Clearly this instance may be replicated to produce an arbitrarily large SR instance for which the size of a maximum Pareto optimal matching is twice the size of a matching output by Tan's algorithm.

Secondly, a related property to Pareto optimality that has been studied is so-called *exchange-stability*. A matching M in an SRC instance is *exchange-stable* [2] if there are no two agents a_i, a_j, each of whom prefers the other's partner to his own partner. Exchange-stability and Pareto optimality are distinct concepts: one may construct instances and matchings to show that each property need not imply the other. For example, consider the two SRC instances I_1 and I_2 as shown in Figure 1 (in an agent's preference list, the symbol . . . denotes all remaining agents listed in arbitrary strict order). In I_1, the matching $M_1 = \{\{1,3\},\{2,4\}\}$ is Pareto

optimal but not exchange-stable. In I_2, the matching $M_2 = \{\{1,2\}, \{3,4\}, \{5,6\}\}$ is exchange-stable but not Pareto optimal. Moreover an SRC instance I need not admit an exchange-stable matching [2], and the problem of deciding whether I does admit such a matching is NP-complete [3, 4].

2 Preliminary definitions and results

We begin this section by defining a property of a matching M that forms a necessary and sufficient condition for M to be Pareto optimal in an SR instance I. In what follows, $bp_I(M)$ denotes the set of blocking pairs with respect to M in I (we omit the subscript if the instance is clear from the context).

Definition 1 *Let M be a matching in an instance of SR. An* improving coalition *with respect to M is a sequence of distinct agents $C = \langle a_0, a_1, \ldots, a_{2r-1}\rangle$, for some $r \geq 1$, such that:*

1. $\{a_{2i-1}, a_{2i}\} \in M$ (1 \leq i \leq r-1);

2. $\{a_{2i}, a_{2i+1}\} \in bp(M)$ (0 \leq i \leq r-1);

3. Either (a) a_0, a_{2r-1} are unmatched in M, or (b) $r \geq 2$ and $\{a_0, a_{2r-1}\} \in M$.

If C satisfies Condition 3(a), we also refer to C as an augmenting coalition, *otherwise we also refer to C as a* cyclic coalition. *Henceforth all subscripts are taken module $2r$ in the context of improving coalitions. We define the* size *of C to be $2r$.*

If M admits no improving (resp. augmenting, cyclic) coalition, we say that M is improving *(resp.* augmenting, cylic*) coalition-free. The matching*

$$M' = (M \backslash \{\{a_{2i-1}, a_{2i}\} : 1 \leq i \leq r\}) \cup \{\{a_{2i}, a_{2i+1}\} : 0 \leq i \leq r-1\}$$

is defined to be the matching obtained from M by satisfying *C. (We remark that if C is an augmenting coalition then $\{a_0, a_{2r-1}\} \notin M$.)*

A matching M is *maximal* in G if $M \cup \{e\}$ is not a matching for any $e \in E \backslash M$. By Definition 1, M is maximal if and only if M admits no improving coalition of size 2. The following proposition indicates that Pareto optimality is equivalent to the absence of an improving coalition.

Proposition 2 *Let M be a matching in a given instance I of SR. Then M is Pareto optimal if and only if M is improving coalition-free.*

Proof: Let M be a Pareto optimal matching in I. If M admits an improving coalition C, let M' be the matching obtained by satisfying C. Then $M' \prec M$, a contradiction.

Conversely let M be a matching that is improving coalition-free, and suppose for a contradiction that M is not Pareto optimal. Then there exists some matching M' such that $M' \prec M$. Let $H = M \oplus M'$ (i.e. H is the graph obtained by taking the symmetric difference of M and M') and let C be a connected component of H. Then C is a path or cycle whose edges alternate between M and M'. We consider two cases.

```
set each agent to be unlabelled;
M := ∅;
for each agent aᵢ
        if aᵢ is unlabelled
                if aᵢ finds some unlabelled agent acceptable
                        let aⱼ be the most-preferred such agent;
                        label each of aᵢ and aⱼ;
                        M := M ∪ {{aᵢ, aⱼ}};
```

Figure 2: Algorithm Greedy-POM.

- *Case (i):* C is an alternating path. As $M' \prec M$, it follows that both end-edges of C belong to M'. Moreover as $M' \prec M$, each agent a in C who is matched in M is also matched in M', and a prefers M' to M. Hence C is an augmenting coalition with respect to M, a contradiction.

- *Case (ii):* C is an alternating cycle. Each agent a in C is matched in both M and M', and as in Case (ii), a prefers M' to M. Hence C is a cyclic coalition with respect to M, a contradiction.

Hence M is Pareto optimal in I. ∎

We now show that Proposition 2 leads to an $O(m)$ algorithm for checking a matching for Pareto optimality in an instance I of SR. Let M be a matching in I and let G be the underlying graph of I. We form a subgraph G_M of G by letting G_M contain only those edges that belong to $M \cup bp(M)$; any isolated vertices are removed from G_M. By Proposition 2, M is Pareto optimal in I if and only if M admits no augmenting path or alternating cycle in G_M. We may test for the existence of the former structure in $O(m)$ time [5, 7]. For the latter structure, we remove any unmatched vertices from G_M (and any edges incident to them) and apply the $O(m)$ alternating cycle detection algorithm of [6]. This discussion leads to the following conclusion.

Proposition 3 *Let M be a matching in a given instance of SR. Then we may check whether M is Pareto optimal in $O(m)$ time.*

We next note that every instance of SR admits at least one Pareto optimal matching, and such a matching may be found in $O(m)$ time using Algorithm Greedy-POM as shown in Figure 2. The correctness and complexity of this algorithm is established by the following proposition.

Proposition 4 *Let I be an instance of SR. Then Algorithm Greedy-POM finds a Pareto optimal matching in I in $O(m)$ time.*

Proof: Without loss of generality suppose that the greedy algorithm considered the agents in order $a_1, a_2, \ldots, a_n$ and constructed the matching M. Now suppose that $M' \prec M$ for some matching M'. Let i be the smallest integer ($1 \leq i \leq n$) such that a_i prefers M' to M. Then a_i is matched in M', to a_j say. When a_i was considered by the algorithm, either (i) a_i was already labelled, or (ii) a_j was already labelled. In Case (i), a_i became matched in M to some a_r when the algorithm considered a_r, where $r < i$. In Case (ii), a_j became matched in M to some a_k when the algorithm

considered either a_j or a_k – let $r = \min\{j, k\}$; then $r < i$. In both cases a_r does not have the same partner in M' as in M. But $M' \prec M$, so a_r prefers $M'(a_r)$ to $M(a_r)$, contradicting the choice of i. Clearly the running time is bounded by the total length of the preference lists. ∎

We now show that stability is a stronger condition than Pareto optimality.

Proposition 5 *Let I be an instance of SR and let M be a matching in I. Then M is stable implies that M is Pareto optimal.*

Proof: By Definition 1, if there is an improving coalition with respect to M then $bp(M) \neq \emptyset$. But M is stable, so that M is Pareto optimal by Proposition 2. ∎

It is easy to construct an SR instance I that admits Pareto optimal matchings of different sizes. Suppose there are four agents, a_1, a_2, a_3, a_4, where $A_1 = \{a_2, a_4\}$, $A_2 = \{a_1\}$, $A_3 = \{a_4\}$, $A_4 = \{a_1, a_3\}$, a_1 prefers a_4 to a_2 and a_4 prefers a_1 to a_3. Then $M_1 = \{\{a_1, a_4\}\}$ is stable (and hence Pareto optimal by Proposition 5) and $M_2 = \{\{a_1, a_2\}, \{a_3, a_4\}\}$ is Pareto optimal. Moreover Algorithm Greedy-POM constructs M_1 given the agent ordering $\langle a_1, a_2, a_3, a_4 \rangle$, and constructs M_2 given the agent ordering $\langle a_2, a_1, a_3, a_4 \rangle$. By creating r copies of I, we may construct an SR instance I^r with $4r$ agents that admits a stable matching M_1^r of size r and a Pareto optimal matching M_2^r of size $2r$. By Propositions 5 and 2, each of M_1^r and M_2^r is a maximal matching in the graph G^r underlying I^r. But the sizes of maximal matchings in G^r differ by at most a factor of 2 [16]. Hence we have an infinite family of instances for which the size of a stable matching is the smallest possible compared to the size of a maximum Pareto optimal matching.

3 Maximum Pareto optimal matchings

Given an SR instance I with underlying graph $G = (A, E)$, a maximum Pareto optimal matching in I may be constructed in polynomial time by imposing weights on the edges of G as follows. For each edge $\{a_i, a_j\} \in E$, the weight of this edge is $rank_i(j) + rank_j(i)$ where $rank_i(j)$ denotes the rank of a_j in a_i's preference list. We may construct a minimum weight maximum cardinality matching M in G in $O(\sqrt{n}\alpha(m, n)m \log^{3/2} n)$ time [8]. The following result indicates that M is a maximum Pareto optimal matching.

Proposition 6 *Let M be a minimum weight maximum cardinality matching in the weighted graph G. Then M is a maximum Pareto optimal matching in I.*

Proof: Suppose not. Then $M' \prec M$ for some matching M'. Every agent matched in M is also matched in M', so $|M'| \geq |M|$. But M is a maximum matching in G, so $|M'| = |M|$, and it follows that the same set of agents are matched in M and M' – we denote these agents by A'. Since $M' \prec M$, for any agent $a_i \in A'$ it follows that $rank_i(M'(i)) \leq rank_i(M(i))$. Moreover, there exists some $a_j \in A'$ such that $rank_j(M'(j)) < rank_j(M(j))$. Hence if $wt(M)$ denotes the weight of matching M, we have

$$wt(M') = \sum_{\{a_i, a_j\} \in M'} (rank_i(j) + rank_j(i)) = \sum_{a_i \in A'} rank_i(M'(i)) < \sum_{a_i \in A'} rank_i(M(i)) = wt(M)$$

which is a contradiction. Hence M is Pareto optimal. ■

Note that the above proposition also indicates that the size of a maximum Pareto optimal matching in I is equal to the size of a maximum matching in G. We remark that an arbitrary matching M in G may be transformed into a Pareto optimal matching M' in I, where $M' \prec M$ and $|M'| \geq |M|$, by repeatedly finding and satisfying improving coalitions. By the discussion preceding Proposition 3, we may find and satisfy an improving coalition with respect to M in $O(m)$ time if one exists. These operations may be repeated until no improving coalition is found, which must occur within m iterations, so the overall process takes $O(m^2)$ time. If M is a maximum matching in G, then all improving coalitions are cyclic coalitions, so $|M'| = |M|$ and M' is a maximum Pareto optimal matching. It remains to consider whether a maximum Pareto optimal matching can be constructed in $O(\sqrt{n}m)$ time. This is the complexity of the fastest current algorithm for finding a maximum matching in a general graph [18].

4   Matchings with fewest number of blocking pairs

We begin this section by presenting two useful results concerned with blocking pairs relative to matchings.

Proposition 7 *Let I be an instance of SR and let M, M' be two matchings in I. Then if $M' \prec M$, it follows that $bp(M') \subset bp(M)$.*

Proof: We firstly show that $bp(M') \subseteq bp(M)$. Let $\{a_i, a_j\} \in bp(M')$. Since neither a_i nor a_j prefers M to M', it follows that $\{a_i, a_j\} \in bp(M)$. We now show that there exists some $\{a_i, a_j\} \in bp(M) \backslash bp(M')$. Let $H = M \oplus M'$ and let C be a connected component of H. Then C is a path or cycle whose edges alternate between M and M'. We consider two cases.

- *Case (i):* C is an alternating path. As $M' \prec M$, it follows that both end-edges of C belong to M'. If C is of length 1, the single edge of C belongs to $bp(M) \backslash bp(M')$. Otherwise, let $\{a_i, a_j\}$ be an end-edge of C. Then without loss of generality a_j is matched in M, to a_k say. But then a_j prefers a_i to a_k, since $M' \prec M$. Hence $\{a_i, a_j\} \in bp(M) \backslash bp(M')$.

- *Case (ii):* C is an alternating cycle. Pick any edge of C that belongs to M', say $\{a_i, a_j\}$. Then neither a_i nor a_j has the same partner in M and M'. But $M' \prec M$, so that $\{a_i, a_j\} \in bp(M) \backslash bp(M')$. ■

Corollary 8 *Let I be an instance of SR and let M be a matching with the fewest number of blocking pairs. Then M is Pareto optimal in I.*

Proof: Suppose not. Then some matching M' satisfies $M' \prec M$. By Proposition 7, $|bp(M')| < |bp(M)|$, a contradiction. Hence M is Pareto optimal in I. ■

4.1 NP-hardness proof

We now consider the problem of finding a matching with the fewest number of blocking pairs (which is necessarily Pareto optimal by Corollary 8). Let Min-BP denote the problem deciding, given an instance I of SR and an integer K, whether I admits a matching M such that $|bp(M)| \leq K$. Define also Min-MM (respectively Exact-MM) to be the problem of deciding, given a graph G and integer K, whether G admits a maximal matching of size at most (respectively exactly) K. Note that Min-MM is NP-complete, even for cubic graphs [12]. It also follows that Exact-MM is NP-complete for the same class of graphs, as we now demonstrate.

Lemma 9 *Exact-MM is NP-complete, even for cubic graphs.*

Proof: Clearly Exact-MM belongs to NP. To show NP-hardness, we reduce from Min-MM, which is NP-complete even for cubic graphs [12]. Let G (a cubic graph) and K (a positive integer) be an instance of the latter problem. Without loss of generality we may assume that $K \leq \beta(G)$, where $\beta(G)$ denotes the size of a maximum matching of G. Suppose that G admits a maximal matching M, where $|M| = k \leq K$. If $k = K$, we are done. Otherwise suppose that $k < K$. Let M' be a maximum matching in G. Now let $H = M \oplus M'$. Each connected component of H is either a path or cycle whose edges alternate between M and M'. Since $|M| = k < K \leq |M'|$, it follows that M admits at least $K - k$ disjoint augmenting paths in H, which we denote by $P_1, \ldots, P_{K-k}$. Let $P = P_1 \cup \ldots \cup P_{K-k}$ and let $M'' = M \oplus P$. Then $|M''| = K$ and the maximality of M implies that M'' is also a maximal matching. The converse is clear. ∎

We now prove that Min-BP is NP-complete, even when all preference lists are complete, using a reduction from Exact-MM.

Theorem 10 *Min-BP is NP-complete, even for complete preference lists.*

Proof: Clearly Min-BP belongs to NP. To show NP-hardness, we reduce from Exact-MM restricted to cubic graphs, which is NP-complete by Lemma 9. Let $G = (V, E)$ (a cubic graph) and K (a positive integer) be an instance of Exact-MM. Assume that $V = \{v_1, \ldots, v_n\}$. Create a new graph $G' = (W, E')$ such that $W = \{w_{i,r} : 1 \leq i \leq n \wedge 1 \leq r \leq 3\}$, each vertex in W is of degree 1 in G', and each edge $\{v_i, v_j\} \in E$ corresponds to a unique edge $\{w_{i,r}, w_{j,s}\} \in E'$ $(1 \leq r, s \leq 3)$; denote this latter edge by $c(\{v_i, v_j\})$. For each $w_{i,r} \in W$, let $e(w_{i,r})$ denote the unique $w_{j,s}$ such that $\{w_{i,r}, w_{j,s}\} \in E'$. For each i $(1 \leq i \leq n)$, define $W_i = \{w_{i,1}, w_{i,2}, w_{i,3}\}$ and $E_i = \{w_{j,3} : \{v_i, v_j\} \in E\}$.

Define also $W' = \{w'_{i,r} : w_{i,r} \in W\}$, $W'' = \{w''_{i,r} : w_{i,r} \in W\}$, $H = \{h_i : 1 \leq i \leq n - 2K\}$, $H' = \{h'_i : h_i \in H\}$ and $H'' = \{h''_i : h_i \in H\}$. We create an instance I of SRC in which the agents are $A = W \cup W' \cup W'' \cup H \cup H' \cup H''$, so that $|A| = 6(2n - K)$. The preference lists of the agents in I are given in Figure 4.1 (assume that $1 \leq i \leq n$, $1 \leq j \leq n - 2K$ and $1 \leq r \leq 3$). In a given agent a's preference list, $[S]$ denotes all members of the set S listed in arbitrary strict order at the position in which the symbol appears. Also $\ldots$ denotes all agents other than a who have not been explicitly listed elsewhere on a's list – such agents are listed in arbitrary strict order at the position in which the symbol appears. We remark that, in the case of $w_{i,3}$'s list, it is possible that $e(w_{i,3}) \notin E_i$. For each $w_{i,r} \in W$, any

8

$$\begin{array}{ll}
w_{i,1}: & w_{i,2} \quad w_{i,3} \quad e(w_{i,1}) \quad w'_{i,1} \quad w''_{i,1} \quad \ldots \\
w_{i,2}: & w_{i,3} \quad w_{i,1} \quad e(w_{i,2}) \quad w'_{i,2} \quad w''_{i,2} \quad \ldots \\
w_{i,3}: & w_{i,1} \quad w_{i,2} \quad e(w_{i,3}) \quad [E_i \backslash \{e(w_{i,3})\}] \quad h_1 \quad h_2 \quad \ldots \quad h_{n-2K} \quad w'_{i,3} \quad w''_{i,3} \quad \ldots \\
w'_{i,r}: & w''_{i,r} \quad w_{i,r} \quad \ldots \\
w''_{i,r}: & w_{i,r} \quad w'_{i,r} \quad \ldots \\
h_j: & w_{1,3} \quad w_{2,3} \quad \ldots \quad w_{n,3} \quad h'_j \quad h''_j \quad \ldots \\
h'_j: & h''_j \quad h_j \quad \ldots \\
h''_j: & h_j \quad h'_j \quad \ldots
\end{array}$$

Figure 3: Preference lists in the constructed SR instance I.

agent whom $w_{i,r}$ prefers to $w'_{i,r}$ is defined to be a *proper agent* for $w_{i,r}$. Similarly, for each $h_j \in H$, any agent whom h_j prefers to h'_j is defined to be a *proper agent* for h_j.

We claim that G admits a maximal matching of size K if and only if I admits a matching with at most n blocking pairs.

For, suppose that G admits a maximal matching M, where $|M| = K$. We form a matching M' in I as follows. Suppose that $\{v_i, v_j\} \in M$. Let $\{w_{i,r}, w_{j,s}\} = c(\{v_i, v_j\})$, let $\{r', r''\} = \{1, 2, 3\} \backslash \{r\}$ and let $\{s', s''\} = \{1, 2, 3\} \backslash \{s\}$. Without loss of generality choose r' and s' such that $w_{i,r}$ and $w_{j,s}$ is the first choice of $w_{i,r'}$ and $w_{j,s'}$ respectively. Add $\{w_{i,r}, w_{j,s}\}$, $\{w_{i,r'}, w_{i,r''}\}$ and $\{w_{j,s'}, w_{j,s''}\}$ to M'. These three assignments imply that $\{\{w_{i,r}, w_{i,r'}\}, \{w_{j,s}, w_{j,s'}\}\} \subseteq bp(M')$.

There are $n - 2K$ vertices in G that are unmatched in M. Let $j_1 < j_2 < \ldots < j_{n-2K}$ be an increasing sequence of integers such that v_{j_r} is unmatched in M ($1 \leq r \leq n-2K$). Add $\{w_{j_r,1}, w_{j_r,2}\}$ and $\{w_{j_r,3}, h_r\}$ to M' ($1 \leq r \leq n-2K$). These two assignments imply that $\{w_{j_r,2}, w_{j_r,3}\} \in bp(M')$. Finally, for each $w_{i,r} \in W$, add $\{w'_{i,r}, w''_{i,r}\}$ to M', and for each $h_j \in H$, add $\{h'_j, h''_j\}$ to M'.

We claim that M' is a matching in I such that $|bp(M')| = n$. For, it is clear that no member of $W' \cup W'' \cup H' \cup H''$ can be involved in a blocking pair of M'. Neither can any $h_l \in H$, since if $\{w_{j,3}, h_l\} \in M'$ for some $w_{j,3} \in W$, and h_l prefers some $w_{i,3} \in W$ to $w_{j,3}$, then either $w_{i,3}$ has a partner of rank ≤ 3 in M' (whom he prefers to h_l), or $\{w_{i,3}, h_k\} \in M'$ for some $h_k \in H$. But then $i < j$, so that $k < l$ by construction of M', so $w_{i,3}$ prefers h_k to h_l. Finally suppose that $\{w_{i,3}, h_k\} \in M'$ for some $w_{i,3} \in W$ and $h_k \in H$. Then v_i is unmatched in M, and hence by the maximality of M, each $w_{j,r} \in E_i \cup \{e(w_{i,3})\}$ has a partner of rank ≤ 3 (whom he prefers to $w_{i,3}$) in M'. Thus every blocking pair of M' comprises a pair of agents belonging to some W_i ($1 \leq i \leq n$).

As previously mentioned, every edge of M gives rise to exactly two blocking pairs of M', and furthermore, every vertex in G that is unmatched in M gives rise to exactly one blocking pair of M'. By the above paragraph these are all the blocking pairs of M', and hence $|bp(M')| = 2K + (n - 2K) = n$.

Conversely suppose that M' is a matching in I such that $|bp(M')| = k \leq n$. We firstly show that $k \geq n$. For, let i be given ($1 \leq i \leq n$). If $\{w_{i,1}, w_{i,2}\} \in M'$ then $\{w_{i,2}, w_{i,3}\} \in bp(M')$. If $\{w_{i,1}, w_{i,3}\} \in M'$ then $\{w_{i,1}, w_{i,2}\} \in bp(M')$. Finally if $w_{i,1}$ has a partner of rank ≥ 3 in M' then $\{w_{i,1}, w_{i,3}\} \in bp(M')$. Hence for each i

$(1 \leq i \leq n)$ there exists a blocking pair of M' comprising a pair of agents belonging to W_i, so that $k \geq n$. Moreover, by the previous inequality, $k = n$ and these are all the blocking pairs of M'.

We next claim that each $w_{i,r} \in W$ has a proper agent as his partner in M'. For if $\{w_{i,r}, w'_{i,r}\} \in M'$ then $\{w'_{i,r}, w''_{i,r}\} \in bp(M')$. If $\{w_{i,r}, w''_{i,r}\} \in M'$ then $\{w_{i,r}, w'_{i,r}\} \in bp(M')$. Finally if $w_{i,r}$ has a partner worse than $w''_{i,r}$ in M' then $\{w_{i,r}, w''_{i,r}\} \in bp(M')$. By a similar argument it follows that each $h_j \in H$ has a proper agent as his partner in M'.

Next we show that, for each i $(1 \leq i \leq n)$, at most one member of W_i can have a partner of rank ≥ 3 in M'. For if two members of W_i have a partner of rank ≥ 3 in M' then all members of W_i do, and hence each of $\{w_{i,1}, w_{i,2}\}$, $\{w_{i,1}, w_{i,3}\}$, $\{w_{i,2}, w_{i,3}\}$ belongs to $bp(M')$, so that $|bp(M')| \geq n + 2$, a contradiction. Hence the set $M = \{\{v_i, v_j\} : \{w_{i,r}, w_{j,s}\} \in M'\}$ is a matching in G.

Now each $h_j \in H$ has a partner in M' whom he prefers to h'_j, so there exists a sequence k_r of integers such that $\{w_{k_r,3}, h_r\} \in M'$ $(1 \leq r \leq n - 2K)$. By the preceding argument, $\{w_{k_r,1}, w_{k_r,2}\} \in M'$ $(1 \leq r \leq n - 2K)$, so that $|M| \leq K$. But for each i $(1 \leq i \leq n)$, some $w_{i,r} \in W_i$ has a partner of rank ≥ 3, and since $w_{i,r}$ has a proper agent as his partner in M', it follows that $|M| = K$.

Finally we show that M is maximal in G. For suppose that two vertices $v_i, v_j \in V$ are unmatched in M, where $\{v_i, v_j\} \in E$. Then by construction of M, $\{w_{i,3}, h_k\} \in M'$ and $\{w_{j,3}, h_l\} \in M'$, for some $h_k, h_l \in H$. But then $\{w_{i,3}, w_{j,3}\} \in bp(M')$, a contradiction. ∎

Note that, by Corollary 8, the above reduction also proves that, given an instance of SRC, the problem of finding a *Pareto optimal* matching with the fewest number of blocking pairs is NP-hard.

4.2 Polynomial-time algorithm for fixed K

We now consider the case where I is an SR instance with underlying graph $G = (A, E)$ and $K \geq 1$ is a fixed constant. We give an $O(m^{K+1})$ algorithm that finds a matching M with at most K blocking pairs, or reports that no such matching exists. We further show how to extend this algorithm if M is required to be Pareto optimal and/or of maximum cardinality.

Our algorithm is based on generating subsets B of edges of G, where $|B| \leq K$ – these edges will form the blocking pairs with respect to a matching to be constructed in a subgraph of G. Given such a set B, we form a subgraph $G_B = (A, E_B)$ of G as follows. For each agent a_i incident to an edge $e = \{a_i, a_j\} \in B$, if e is a blocking pair of a matching M, it follows that $\{a_i, a_j\} \notin M$ and a_i cannot be matched in M to an agent whom he prefers to a_j in I. Hence we delete $\{a_i, a_j\}$ from E_B, and also we delete $\{a_i, a_k\}$ from E_B for any a_k such that a_i prefers a_k to a_j in I. If any such edge $\{a_i, a_k\}$ is not in B, then we require that $\{a_i, a_k\}$ is not a blocking pair of a constructed matching M. This can only be achieved if a_k is matched in M to an agent whom he prefers to a_i in I. Hence we invoke $truncate_{a_k}(a_i)$, which represents the operation of deleting $\{a_k, a_l\}$ from E_B, for any a_l such that a_k prefers a_i to a_l in I. Additionally we add a_k to a set P to subsequently check that a_k is matched.

Having completed the construction of G_B, we denote by I_B the SR instance with underlying graph G_B and preference lists obtained by restricting the preferences in

```
for each B ⊆ E such that |B| ≤ K
    E_B := E;      // G_B = (A, E_B) is a subgraph of G
    P := ∅;
    for each agent a_i incident to some {a_i, a_j} ∈ B
        delete {a_i, a_j} from E_B;
        for each agent a_k such that a_i prefers a_k to a_j in I
            delete {a_i, a_k} from E_B;
            if {a_i, a_k} ∉ B
                truncate_{a_k}(a_i);
                P := P ∪ {a_k};
    if there is a stable matching M in I_B
        if every agent in P is matched in M
            output M and halt;
report that no matching with ≤ K blocking pairs exists;
```

Figure 4: Algorithm K-BP

I to E_B. Any matching M in G_B satisfies $B \subseteq bp_I(M)$. To avoid any additional blocking pairs in I, we seek a stable matching in I_B in which all agents in P are matched. We apply Irving's algorithm for SR [11] to I_B – suppose it finds a stable matching M in I_B. If all agents in P are matched, then it follows that $bp_I(M) = B$, and hence $|bp(M)| \leq K$ – thus we may output M and halt. If some agents in P are unmatched in M then we need not consider any other stable matching in I_B, since Theorem 4.5.2 of [11] asserts that the same agents are matched in all stable matchings in I_B. Hence (and also in the case that no stable matching in I_B is found), we may consider the next subset B. If we complete the generation of all subsets B without having output a matching M, we report that no matching with the desired property exists. The algorithm is displayed as Algorithm K-BP in Figure 4.

We now establish the correctness and complexity of Algorithm K-BP.

Theorem 11 *Given an SR instance I and a fixed constant K, Algorithm K-BP finds a matching with at most K blocking pairs, or reports that no such matching exists, in $O(m^{K+1})$ time.*

Proof: Suppose firstly that the algorithm outputs a matching M when the outermost loop considered a set B. We show that M is a matching in I such that $bp_I(M) = B$. For, let $\{a_i, a_j\} \in B$. Then by construction of G_B, $\{a_i, a_j\} \notin M$. Moreover either a_i is unmatched in M, or a_i is matched in M and prefers a_j to $M(a_i)$ in I. Similarly either a_j is unmatched in M, or a_j is matched in M and prefers a_i to $M(a_j)$ in I. Hence $\{a_i, a_j\} \in bp_I(M)$, so that $B \subseteq bp_I(M)$. We now show that $bp_I(M) \subseteq B$. For, suppose that $\{a_k, a_l\} \in (E \backslash B) \cap bp_I(M)$. Then $\{a_k, a_l\} \notin E_B$, as M is stable in I_B. Hence $\{a_k, a_l\}$ has been deleted by the algorithm. Thus without loss of generality $a_k \in P$, so that a_k is matched in M and a_k prefers $M(a_k)$ to a_l in I. Hence $\{a_k, a_l\} \notin bp_I(M)$ after all, so that $bp_I(M) = B$.

Now suppose that M is a matching in I such that $bp_I(M) = B$, where $|B| \leq K$. By the above paragraph, if, before considering B, the outermost loop had already output a matching M' when considering a subset B', then $bp_I(M') = B'$, and $|B'| \leq K$. Otherwise, when the outermost loop considers the subset B, it must

be the case that no edge of M is deleted when constructing G_B. Hence $M \subseteq E_B$. Moreover M is stable in I_B, for if not then $e \in bp_{I_B}(M)$ for some $e \in E_B$, and hence $e \in bp_I(M)$. Since $B \cap E_B = \emptyset$, it follows that $e \in bp_I(M) \backslash B$, a contradiction. Finally every member of P is matched in M, for suppose $a_k \in P$ is unmatched in M. As $a_k \in P$, there is some agent a_i such that a_i prefers a_k to a_j in I, where $\{a_i, a_j\} \in B$ and $\{a_i, a_k\} \notin B$. Hence $\{a_i, a_k\} \in bp_I(M) \backslash B$, a contradiction. Hence by [11, Theorem 4.5.2], Irving's algorithm finds a stable matching M' in I_B (possibly $M' = M$) such that all members of P are matched in M'. Thus the algorithm outputs M' in this case. By the above paragraph, $bp_I(M') = B$.

Finally suppose that there is no matching M in I such that $|bp(M)| \le K$. By the first paragraph, if the algorithm outputs a matching M' when the outermost loop considered a subset B, then $bp_I(M') = B$, a contradiction. Hence the algorithm reports that no such matching M exists in this case.

Clearly the outermost loop iterates $O(m^K)$ times. Within a loop iteration, construction of G_B takes $O(m)$ time, as does the invocation of Irving's algorithm. All other operations are $O(m)$, and hence the time complexity of Algorithm K-BP follows. ∎

Note that, if Algorithm K-BP outputs a matching M with at most K blocking pairs, by the discussion at the end of Section 3, we may transform M into a Pareto optimal matching M' such that $M' \prec M$ (and $|M'| \ge |M|$) in $O(m^2)$ time. By Proposition 7, $|bp(M')| \le K$. Also, it is straightforward to modify Algorithm K-BP so that it outputs the largest stable matching taken over all subsets B – we may then find a matching M such that (i) M is Pareto optimal, (ii) $|bp(M)| \le K$, and (iii) M is of maximum cardinality with respect to (i) and (ii). This extension uses the fact that all stable matchings in I_B have the same size [11, Theorem 4.5.2], so that the choice of stable matching constructed by the algorithm is not of significance as far as Condition (iii) is concerned.

5   Minimum Pareto optimal matchings

In this section, we consider minimum Pareto optimal matchings. Let Min-POM denote the problem deciding, given an instance I of SR and an integer K, whether I admits a Pareto optimal matching of size at most K. Define also the problem X3C as follows. An instance $\langle X, C \rangle$ of X3C comprises a set $X = \{x_1, x_2, \ldots, x_n\}$ of *elements*, where $n = 3q$ for some integer q, and a collection $C = \{c_1, c_2, \ldots, c_m\}$ of subsets of X (*clauses*), where $|c_i| = 3$ for each $c_i \in C$. The problem is to decide whether C contains an *exact cover* for X, i.e. whether there exist a set $C' \subseteq C$ of pairwise disjoint sets whose union is X. X3C is NP-complete [10, problem SP2].

We now prove that Min-POM is NP-complete via a reduction from X3C. The following proof is adapted from a construction of McRae [17].

Theorem 12 *Min-POM is NP-complete.*

Proof: By Proposition 3, Min-POM belongs to NP . To show NP-hardness, we reduce from X3C. Hence let $\langle X, C \rangle$ be an instance of X3C as defined above. Let i_r $(1 \le r \le 3m)$ be a sequence of integers such that, for each j $(1 \le j \le m)$, $c_j = \{x_{i_{3j-2}}, x_{i_{3j-1}}, x_{i_{3j}}\}$. We construct an instance I of SR as follows. Let $W =$

$\{w_1, w_2, \ldots, w_{3q}\}$, let $Y = \{y_1, y_2, \ldots, y_{3m}\}$ and let $Z = \{z_1, z_2, \ldots, z_m\}$. The agents in J are $W \cup X \cup Y \cup Z$. Let $F = \{\{w_i, x_i\} : 1 \leq i \leq 3q\}$. Also, for each i $(1 \leq i \leq 3q)$, define $Y_i = \{y_j \in Y : c_j \in C \wedge x_i \in c_j\}$. We now create preference lists for the agents in I as follows (assume that $1 \leq i \leq 3q$, $1 \leq j \leq m$ and $0 \leq k \leq 2$):

$$
\begin{array}{ll}
w_i : & x_i \\
x_i : & [Y_i] \quad w_i \\
y_{3j-k} : & x_{i_{3j-k}} \quad z_j \\
z_j : & y_{3j-2} \quad y_{3j-1} \quad y_{3j}
\end{array}
$$

Let $K = m + 2q$. We claim that $\langle X, C \rangle$ has an exact cover if and only if I admits a Pareto optimal matching of size at most K.

For, suppose that C' is an exact cover for X. Then $|C'| = q$. We define a matching M in I as follows. For each $c_j \in C'$, add $\{y_{3j-k}, x_{i_{3j-k}}\}$ to M $(0 \leq k \leq 2)$. For each $c_j \notin C'$, add $\{z_j, y_{3j-1}\}$ to M. Then $|M| = 3q + (m - q) = K$. It is straightforward to verify that M is maximal in I. Moreover M admits no improving coalition. For, any augmenting path or alternating cycle relative to M must involve an agent $x_i \in X$. But $\{x_i, y_j\} \in M$ for some $y_j \in Y_i$, and x_i is the first choice of y_j. Hence M is Pareto optimal, as required.

Conversely suppose that I admits a Pareto optimal matching of size at most K. Choose M to be such a matching that minimises $|M \cap F|$. We claim that $M \cap F = \emptyset$. For, suppose $\{w_i, x_i\} \in M \cap F$. Pick any $y_l \in Y_i$ and suppose that $y_l = y_{3j-k}$ $(1 \leq j \leq m, 0 \leq k \leq 2)$. We consider three cases:

- *Case 1*: $\{z_j, y_{3j-k}\} \notin M$. Let $M' = (M \backslash \{\{w_i, x_i\}\}) \cup \{\{x_i, y_{3j-k}\}\}$.

- *Case 2*: $\{z_j, y_{3j-k}\} \in M$ and y_{3j-r} is unmatched in M for some r $(0 \leq r \leq 3)$. Let $M' = (M \backslash \{\{w_i, x_i\}, \{z_j, y_{3j-k}\}\}) \cup \{\{x_i, y_{3j-k}\}, \{z_j, y_{3j-r}\}\}$.

- *Case 3*: $\{z_j, y_{3j-k}\} \in M$ and y_{3j-r} is matched in M for each $r \in \{0, 1, 2\} \backslash \{k\}$. Let $M' = (M \backslash \{\{w_i, x_i\}, \{z_j, y_{3j-k}\}\}) \cup \{\{x_i, y_{3j-k}\}\}$.

Then in each case, it may be verified that $|M'| \leq |M|$ and the Pareto optimality of M implies that M' is Pareto optimal. Moreover $|M' \cap F| < |M \cap F|$, contradicting the choice of M. Hence the claim is established.

For each i $(1 \leq i \leq 3q)$, $\{x_i, y_j\} \in M$ for some $y_j \in Y_i$, for otherwise M is not maximal. Let p denote the number of edges of M incident to vertices in Z. Then $|M| \geq 3q + p$. Hence $p \leq m - q$. Define a vertex $z_j \in Z$ to be *saturated* by M if $\{y_{3j-r}, x_{i_{3j-r}}\} \in M$ for $0 \leq r \leq 2$. Let t denote the number of vertices in Z that are saturated by M. Then $3t \leq 3q$. Also it follows that $p = m - t$, for $z_j \in Z$ is matched in M if and only if z_j is not saturated by M, by maximality of M. Hence $t = q$, so that $C' = \{c_j \in C : z_j \text{ is saturated by } M\}$ is an exact cover for X. ∎

It is clear that the underlying graph corresponding to the SR instance as constructed by the above reduction is bipartite. Furthermore, X3C remains NP-complete even if each member of X occurs in either two or three clauses [10, problem SP2]. Hence Min-POM remains NP-complete even for instances of the Stable Marriage problem with incomplete preference lists in which each agent's list is of length at most 4.

As observed at the end of Section 2, the size of a minimum Pareto optimal matching is at least half the size of a maximum Pareto optimal matching. Hence Algorithm Greedy-POM is a 2-approximation algorithm for the problem of finding a minimum Pareto optimal matching.

6 Concluding remarks

The results of this paper leave open some interesting questions. Firstly as mentioned in Section 3, it remains to consider whether there is a faster algorithm for finding a maximum Pareto optimal matching in an SR instance. Secondly, the results of Section 4 leave open the question of the approximability of the problem of finding a matching with the minimum number of blocking pairs. Finally, the case where preference lists in SR may include ties merits further investigation.

Acknowledgement

References

[1] D.J. Abraham, K. Cechlárová, D.F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. To appear in *Proceedings of ISAAC 2004: the 15th Annual International Symposium on Algorithms and Computation.* Lecture Notes in Computer Science, Springer-Verlag, 2004.

[2] J. Alcalde. Exchance-proofness or divorce-proofness? Stability in one-sided matching markets. *Economic Design*, 1:275–287, 1995.

[3] K. Cechlárová. On the complexity of exchange-stable roommates. *Discrete Applied Mathematics*, 116(3):279–287, 2002.

[4] K. Cechlárová and D.F. Manlove. The Exchange-stable Marriage Problem. Technical Report TR-2003-142, University of Glasgow, Department of Computing Science, September 2003.

[5] H.N. Gabow. An efficient implementations of Edmonds' algorithm for maximum matching on graphs. *Journal of the ACM*, 23(2):221–234, 1976.

[6] H.N. Gabow, H. Kaplan, and R.E. Tarjan. Unique maximum matching algorithms. *Journal of Algorithms*, 40:159–183, 2001.

[7] H.N. Gabow and R.E. Tarjan. A linear-time algorithm for a special case of disjoint set union. *Journal of Computer and System Sciences*, 30:209–221, 1985.

[8] H.N. Gabow and R.E. Tarjan. Faster scaling algorithms for general graph-matching problems. *Journal of the ACM*, 38(4):815–853, 1991.

[9] D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

[10] M.R. Garey and D.S. Johnson. *Computers and Intractability.* Freeman, San Francisco, CA., 1979.

[11] D. Gusfield and R.W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.

[12] J.D. Horton and K. Kilakos. Minimum edge dominating sets. *SIAM Journal on Discrete Mathematics*, 6:375–387, 1993.

[13] R.W. Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6:577–595, 1985.

[14] R.W. Irving and D.F. Manlove. The Stable Roommates Problem with Ties. *Journal of Algorithms*, 43:85–105, 2002.

[15] D.E. Knuth. *Stable Marriage and its Relation to Other Combinatorial Problems*, volume 10 of *CRM Proceedings and Lecture Notes*. American Mathematical Society, 1997. English translation of *Mariages Stables*, Les Presses de L'Université de Montréal, 1976.

[16] B. Korte and D. Hausmann. An analysis of the greedy heuristic for independence systems. In *Annals of Discrete Mathematics*, volume 2, pages 65–74. North-Holland, 1978.

[17] A.A. McRae. *Generalizing NP-completeness proofs for bipartite and chordal graphs*. PhD thesis, Clemson University, Department of Computer Science, South Carolina, 1994.

[18] S. Micali and V.V. Vazirani. An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In *Proceedings of FOCS '80: the 21st Annual IEEE Symposium on Foundations of Computer Science*, pages 17–27. IEEE Computer Society, 1980.

[19] B.G. Pittel and R.W. Irving. An upper bound for the solvability probability of a random stable roommates instance. *Random Structures and Algorithms*, 5(3):465–486, 1994.

[20] A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis*, volume 18 of *Econometric Society Monographs*. Cambridge University Press, 1990.

[21] J.J.M. Tan. Stable matchings and stable partitions. *International Journal of Computer Mathematics*, 39:11–20, 1991.

[22] A. Thomson. An empirical study of the stable roommates problem. BSc Honours project report, University of Glasgow, Department of Computing Science, 2004.