The Stable Roommates Problem with Globally-Ranked Pairs

David J. Abraham[1]*, Ariel Levavi[1],
David F. Manlove[2]**, and Gregg O'Malley[2]

[1] Computer Science Department, Carnegie Mellon University, USA.
dabraham@cs.cmu.edu, alevavi@andrew.cmu.edu
[2] Department of Computing Science, University of Glasgow, UK.
davidm@dcs.gla.ac.uk, gregg@dcs.gla.ac.uk

Abstract. We introduce a restriction of the stable roommates problem in which roommate pairs are ranked globally. In contrast to the unrestricted problem, weakly stable matchings are guaranteed to exist, and additionally, can be found in polynomial time. However, it is still the case that strongly stable matchings may not exist, and so we consider the complexity of finding weakly stable matchings with various desirable properties. In particular, we present a polynomial-time algorithm to find a rank-maximal (weakly stable) matching. This is the first generalization of the algorithm due to Irving et al. [18] to a non-bipartite setting. Also, we prove several hardness results in an even more restricted setting for each of the problems of finding weakly stable matchings that are of maximum size, are egalitarian, have minimum regret, and admit the minimum number of weakly blocking pairs.

1 Introduction

The STABLE ROOMMATES problem (SR) [11, 16, 15, 17] involves pairing-up a set of *agents*, each of whom ranks the others in (not necessarily strict) order of preference. Agents can declare each other *unacceptable*, in which case they cannot be paired together. Our task is to find a pairing of mutually acceptable agents such that no two agents would prefer to partner each other over those that we prescribed for them.

We represent acceptable pairs by a graph $G = (V, E)$, with one vertex $u \in V$ for each agent, and an edge $\{u, v\} \in E$ whenever agents u and v are mutually acceptable. A pairing is just a *matching* M of G, i.e. a subset of edges in E, no two of which share a vertex. If $\{u, v\} \in M$, we say that u is *matched* in M and $M(u)$ denotes v, otherwise u is *unmatched* in M. An agent u prefers one matching M' over another M if i) u is matched in M' and unmatched in M, or ii) u prefers $M'(u)$ to $M(u)$. Similarly, u is indifferent between M' and M if i) u is unmatched in M' and M, or ii) u is indifferent between $M'(u)$ and $M(u)$.

* Research supported in part by NSF grants IIS-0427858 and CCF-0514922IIS-0427858. Part of this work completed while visiting Microsoft Research, Redmond.
** Supported by EPSRC grant EP/E011993/1.

A matching M is *weakly stable* if it admits no *strongly blocking pair*, which is an edge $\{u, v\} \in E \backslash M$ such that u and v prefer $\{\{u, v\}\}$ to M. A matching M is *strongly stable* if it admits no *weakly blocking pair*, which is an edge $\{u, v\} \in E \backslash M$ such that u prefers $\{\{u, v\}\}$ to M, while v either prefers $\{\{u, v\}\}$ to M, or is indifferent between them.

In this paper, we introduce and study the STABLE ROOMMATES WITH GLOBALLY-RANKED PAIRS problem (SR-GRP). An instance of SR-GRP is a restriction of SR in which preferences may be derived from a ranking function $rank : E \to \mathbb{N}$. An agent u prefers v to w if $e = \{u, v\}$, $e' = \{u, w\}$ and $rank(e) < rank(e')$, and u is indifferent between them if $rank(e) = rank(e')$.

Before giving our motivation for studying this restriction, we introduce some notation. We define E_i to be the set of edges with rank i, and $E_{\leq i}$ to be the set $E_1 \cup E_2 \cup \ldots \cup E_i$. Additionally, let $n = |V|$ be the number of agents, $m = |E|$ be the number of mutually acceptable pairs. Without loss of generality, we assume the maximum edge rank is at most m. Also, we make the standard assumption in stable marriage problems that the adjacency list for a vertex is given in order of preference/rank.

Motivation. In several real-world settings, agents have restricted preferences that can be represented by the SR-GRP model. A pairwise kidney exchange market [26, 25, 1] is one such setting. Here, patients with terminal kidney-disease obtain compatible donors by swapping their own willing but incompatible donors. We can model the basic market by constructing one vertex for each patient, and an undirected edge between any two patients where the incompatible donor for one patient is compatible with the other patient, and vice versa. Of course, patients may have different preferences over donors. However, since the expected years of life gained from a transplant is similar amongst all compatible kidneys, the medical community has suggested that patient preferences should be *binary/dichotomous* [14, 7] – i.e., patients are indifferent between all compatible donors. Binary preferences are easily modelled in SR-GRP by giving all edges the same rank.

A second example also comes from pairwise kidney exchange markets. When two (patient,donor) pairs are matched with each other (in order to swap donors), we are not certain if the swap can occur until expensive last-minute compatibility tests are performed on the donors and patients. If either potential transplant in the swap is incompatible, the swap is cancelled and the two patients must wait for a future match run. Since doctors can rank potential swaps by their chance of success, and patients prefer swaps with better chances of success, this generalizes the binary preference model above, and can clearly still be modelled by SR-GRP.

One final real-world setting is described in [4]. When colleges pair-up freshmen roommates, it is not feasible for students to rank each other explicitly. Instead, each student submits a form which describes him/herself in several different dimensions (e.g. bedtime preference, cleanliness preference etc). Students can then be represented as points in a multidimensional space, and preferences over other students can be inferred by a distance function. Note that this model

[4] is a restriction of SR-GRP in that it is not possible to declare another student unacceptable.

In order to highlight the generality of the SR-GRP model, we introduce a second restriction of SR called STABLE ROOMMATES WITH GLOBALLY-ACYCLIC PREFERENCES (SR-GAP). Instances of SR-GAP satisfy the following characterization test: given an arbitrary instance I of SR with $G = (V, E)$, construct a digraph $P(G)$, containing one vertex e for each edge in $e \in E$, and an arc from $e = \{u, v\} \in E$ to $e' = \{u, w\} \in E$ if u prefers w to v. Now, for each $e = \{u, v\}$ and $e' = \{u, w\}$ in E, if u is indifferent between v and w, merge vertices e and e'. Note that a merged vertex may contain several original edge-vertices and have self-loops. Instance I belongs to SR-GAP iff $P(G)$ is acyclic.

Instances of SR-GRP satisfy the SR-GAP test, since any directed path in $P(G)$ consists of arcs with monotonically improving ranks, and so no cycles are possible. In the reverse direction, given any instance of SR-GAP, we can derive a suitable rank function from a reverse topological sort on $P(G)$, i.e. $rank(e) < rank(e')$ iff e appears before e'. The following proposition is clear:

Proposition 1. *Let I be an instance of* SR. *Then I is an instance of* SR-GRP *if and only if I is an instance of* SR-GAP.

As well as modelling real-world problems, SR-GRP is an important theoretical restriction of SR. It is well-known that SR has two key undesirable properties. First, some instances of SR admit no weakly stable matchings (see, for example, [15, page 164]). And second, the problem of finding a weakly stable matching, or proving that no such matching exists, is NP-hard [24, 17]. It turns out that SR-GRP has neither of these undesirable properties [4] [3].

Lemma 1. *Let $G = (V, E_1 \cup \ldots \cup E_m)$ be an instance of* SR-GRP. *Then M is a weakly stable matching of G if and only if $M \cap E_{\leq i}$ is a maximal matching of $E_{\leq i}$, for all i.*

So we can construct a weakly stable matching in $O(n + m)$ time by finding a maximal matching on rank-1 edges, removing the matched vertices, finding a maximal matching on rank-2 edges, and so on.

Strongly stable matchings are also easy to characterize in SR-GRP [4].

Lemma 2. *Let $G = (V, E_1 \cup \ldots \cup E_m)$ be an instance of* SR-GRP. *Then M is a strongly stable matching of G if and only if $M \cap E_i$ is a perfect matching of $\{e \in E_i : e$ is not adjacent to any $e' \in E_{<i}\}$, for all i.*

Of course, even E_1 may not admit a perfect matching, and so strongly stable matchings may not exist. However, we can find a strongly stable matching, or prove that no such matching exists in $O(m\sqrt{n})$ time by using the maximum matching algorithm of Micali and Vazirani for non-bipartite graphs [23]. This improves on the best known running time of $O(m^2)$ for general SR [27].

[3] Lemmas 1 and 2 are proved by [4] in a restricted setting. However, their extensions to SR-GRP are straightforward.

These observations show that SR-GRP can be far simpler than SR. In this paper, we are interested in which problems become more tractable in SR-GRP, and which problems maintain their hardness. Work along these lines has been done before [5, 28, 6, 4]. For example, Chung [6] shows that the "no odd ring" condition on preferences is sufficient for the existence of a weakly stable matching. The SR-GAP acyclic condition is a restriction of the "no odd ring" condition, in that neither odd nor even rings are permitted.

The possible non-existence of a strongly stable matching motivates the search for weakly stable matchings with desirable properties. A *rank-maximal* matching [18, 29] includes the maximum possible number of rank-1 edges, and subject to this, the maximum possible number of rank-2 edges, and so on. More formally, define the *signature* of a matching M as $\langle s_1, s_2, \ldots, s_m\rangle$, where s_i is the number of rank-i edges in M. Then a matching is rank-maximal iff it has the lexicographic-maximal signature amongst all matchings.

Recall from Lemma 2 that a strongly stable matching is perfect on rank-1 edges, and subject to this, perfect on rank-2 edges, and so on. It is clear that a rank-maximal matching is strongly stable, when strong stability is possible. If no strongly stable matching exists, then a rank-maximal matching, which by Lemma 1 is always weakly stable, seems a natural substitute. Irving et al. [18] gave an $O(\min(n + R, R\sqrt{n})m)$ algorithm for the problem of finding a rank-maximal matching in a bipartite graph, where R is the rank of the worst-ranked edge in the matching.

Other desirable types of weakly stable matchings may be those that have maximum cardinality, are *egalitarian*, are of *minimum regret*, or admit the fewest number of weakly blocking pairs. An egalitarian (respectively minimum regret) weakly stable matching satisfies the property that the sum of the ranks (respectively the maximum rank) of the edges is minimised, taken over all weakly stable matchings. Given a general SR instance I, each of the problems of finding an egalitarian and a minimum regret weakly stable matching is NP-hard [9, 20] (in the former case, even if the preference lists are complete and strictly-ordered, and in the latter case, even if the underlying graph is bipartite). However the complexity of the problem of finding a weakly stable matching with the minimum number of weakly blocking pairs in I has, until now, been open.

Paper outline and summary of contribution. In Section 2, we consider rank-maximal matchings, and present the first generalization of Irving et al.'s [18] algorithm to a non-bipartite setting. In Section 3, we prove hardness results for for each of the problems of finding weakly stable matchings that are of maximum size, are egalitarian, have minimum regret, and admit the minimum number of weakly blocking pairs. We also show that this last problem is inapproximable within a factor of $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless P = NP. These hardness results apply even in a restricted version of SR-GRP in which the graph G is bipartite, and (in the first three cases) if an agent v is incident to an edge of rank k, then v is incident to an edge of rank k', for $1 \leq k' \leq k$. Finally, Section 4 contains concluding remarks.

2 Rank-Maximal Matching

One obvious way to construct a rank-maximal matching is to find a maximum-weight matching using edge weights that increase exponentially with improving rank. However, with K distinct rank values, Gabow and Tarjan's matching algorithm [10] takes $O(K^2\sqrt{n\alpha(m,n)\lg nm}\lg n)$ time[4], where α is the inverse Ackermann function. As in the bipartite restriction [18], our combinatorial algorithm avoids the problem of exponential-sized edge weights, leading to an improved runtime of $O(\min\{n+R, R\sqrt{n}\}m)$, where $R \leq K$ is the rank of the worst-ranked edge in the matching.

Let $G_i = (V, E_{\leq i})$. Our algorithm begins by constructing a maximum matching M_1 on G_1. Note that M_1 is rank-maximal on G_1 by definition. Then inductively, given a rank-maximal matching M_{i-1} on G_{i-1}, the algorithm exhaustively augments M_{i-1} with edges from E_i to construct a rank-maximal matching M_i on G_i. In order to ensure rank-maximality, certain types of edges are deleted before augmenting. With these edges deleted, it becomes possible to augment M_{i-1} *arbitrarily*, while still guaranteeing rank-maximality. Hence, we can perform the augmentations using Micali and Vazirani's fast maximum matching algorithm [23]. In the non-bipartite setting, we perform one additional type of edge deletion beyond the bipartite setting. Additionally, we shrink certain components into *supervertices*. Note that this shrinking is separate from any blossom-shrinking [8] that might occur in the maximum matching subroutine.

In order to understand the edge deletions and component shrinking, recall the Gallai-Edmonds decomposition lemma [19]: Let $G = (V, E)$ be an arbitrary undirected graph. Then V can be partitioned into the following three sets, namely GED-U$[G]$, GED-O$[G]$ and GED-P$[G]$. Vertices in GED-U$[G]$ are *underdemanded*, since they are unmatched in some maximum matching of G. All other vertices that are adjacent to one in GED-U$[G]$ are *overdemanded* and belong to GED-O$[G]$. Finally, all remaining vertices are *perfectly demanded* and belong to GED-P$[G]$. The decomposition lemma gives many useful structural properties of maximum matchings. For example, in every maximum matching, vertices in GED-O$[G]$ are always matched, and their partner is in GED-U$[G]$. Similarly, vertices in GED-P$[G]$ are always matched, though their partners are also in GED-P$[G]$. We will use the properties given in Lemma 3.

Lemma 3 (Gallai-Edmonds Decomposition). *In any maximum matching M of G,*

1. *For all u in* GED-O$[G]$*, $M(u)$ is in* GED-U$[G]$
2. *For all even (cardinality) components C of $G \setminus$* GED-O$[G]$*, i) $C \subseteq$* GED-P$[G]$*, and ii) $M(u)$ is in C, for all u in C*
3. *For all odd (cardinality) components C of $G \setminus$* GED-O$[G]$*, i) $C \subseteq$* GED-U$[G]$*, ii) $M(u)$ is in C, for all u in C except one, say v, and iii) either v is unmatched in M, or $M(v)$ is in* GED-O$[G]$

[4] See [22] for an explanation of the K^2 factor.

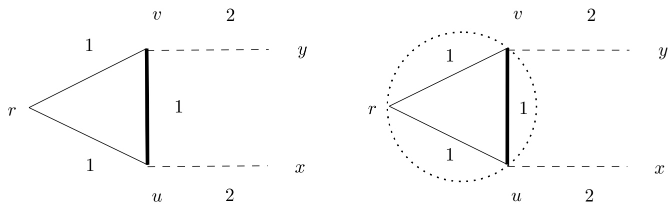

Fig. 1. Example of shrinking operation

Consider the first inductive step of the algorithm, in which we are trying to construct a rank-maximal matching M_2 of $G_2 = (V, E_{\leq 2})$, given a maximum matching M_1 of $G_1 = (V, E_1)$. We do not want to *commit* to edges in M_1 at this point, because perhaps no rank-maximal matching on G_2 contains these edges. However, according to the decomposition lemma, we can safely *delete* any edge $e = \{u, v\}$ such that:

(i) $u \in$ GED-O$[G_1]$ and $v \in$ GED-O$[G_1] \cup$ GED-P$[G_1]$
(ii) $e \in E_{\geq 2}$, and $u \in$ GED-O$[G_1] \cup$ GED-P$[G_1]$
(iii) $e \in E_{\geq 2}$, and both u and v belong to the same odd component of G_1

We delete all such edges to ensure they are not subsequently added to the matching when we augment. Note that the third deletion type is required for non-bipartite graphs, since only one vertex in each odd component C is unmatched internally.

After deleting edges in G_1, we *shrink* each odd component C into a supervertex. We define the *root* r of C as the one vertex in C that is unmatched within C. Note that C's supervertex is matched iff r is matched. Now, when we add in undeleted edges from $e = \{u, v\} \in E_{\geq 2}$ into the graph, if $u \in C$ and $v \notin C$, we replace e with an edge between v and C's supervertex. Note that during the course of the algorithm, we will be dealing with graphs containing supervertices, which themselves, recursively contain supervertices. In such graphs, we define a *legal* matching to be any collection of independent edges such that in every supervertex, all top-level vertices but the root are matched internally.

To give some intuition for why we shrink odd components, consider the graph in Figure 1. The triangle of rank-1 edges is an odd component (with $\{u, v\}$ matched), and so neither rank-2 edges are deleted. One way to augment this graph is to include the two rank-2 edges and take out the rank-1 $\{u, v\}$ edge. This destroys the rank-maximal matching on G_1. If we shrink the triangle however, the supervertex is unmatched, and so $\{x, u\}$ and $\{y, v\}$ are both valid augmenting paths. Note how these augmenting paths can be expanded *inside* the supervertex by removing and adding one rank-1 edge to end at the root r. This expansion makes the augmenting path legal in the original graph, while not changing the number of matching edges internal to the supervertex.

Figure 2 contains pseudocode for our non-bipartite rank-maximal matching algorithm. One aspect that requires more explanation is how we augment M_i in G'_i. The overall approach is to find an augmenting path P while regarding each

Rank-Maximal-Matching$(G = (V, E_1 \cup E_2 \cup \ldots \cup E_m))$
Set G'_1 to G_1;
Let M_1 be any maximum matching of G_1;
For $i = 2$ to m:
Set G'_i to G'_{i-1}, and M_i to M_{i-1};
Compute the GED of G'_{i-1} using M_{i-1};
Delete edges in G'_i between two vertices in GED-O$[G'_{i-1}]$;
Delete edges in G'_i between vertices in GED-O$[G'_{i-1}]$ and GED-P$[G'_{i-1}]$;
Delete any edge e in $E_{\geq i}$ where:
i) e is incident on a GED-O$[G'_{i-1}]$ or GED-P$[G'_{i-1}]$ vertex, or
ii) e is incident on two vertices in the same odd component of G_{i-1};
Shrink each odd component of G_{i-1} in the graph G'_i;
Add undeleted edges in E_i to G'_i;
Augment M_i in G'_i until it is a maximum matching;
End For
Return M_m;

Fig. 2. Non-bipartite rank-maximal matching algorithm

top-level supervertex in G'_i as a regular vertex. Then for each supervertex C in P, we expand P through C in the following way. Let u be the vertex in C that P enters along an unmatched edge. If u is the root r of C, then C is unmatched, and we can replace C by u in P. Otherwise, $u \neq r$, and either C is unmatched or P leaves C via the matched edge incident on r. In the next lemma, we show that there is an even-length alternating path from u to r, beginning with a matched edge. We can expand P by replacing C with this even-length alternating path.

Lemma 4. *Let M be a legal matching on some supervertex C with root r. Let u be any other node in C. Then there is an even-length alternating path from u to r beginning with a matched edge.*

Proof. Let M' be a legal matching of C in which u is unmatched (such a matching is guaranteed by the decomposition lemma). Consider the symmetric difference of M and M'. Since every vertex besides u and r is matched in both matchings, there must be an even-length alternating path consisting of M and M' edges from u to r. □

In all cases of P and C, note that C has the same number of internally matched edges before and after augmentation by P, and so the matching remains legal. Also, if r was matched prior to augmentation, then it is still matched afterwards.

The next three lemmas, which generalize those in [18], establish the correctness of the algorithm. Lemma 5 proves that no rank-maximal matching contains a deleted edge. Lemma 6 proves that augmenting a rank-maximal matching M_{i-1} of G_{i-1} does not change its signature up to rank $(i-1)$. And finally, Lemma 7 proves that the final matching is rank-maximal on the original graph G.

Lemma 5. *Suppose that every rank-maximal matching of G_{i-1} is a maximum legal matching on G'_{i-1}. Then every rank-maximal matching of G_i is contained in G'_i.*

Proof. Let M be an arbitrary rank-maximal matching of G_i. Then $M \cap E_{\leq i-1}$ is a rank-maximal matching of G_{i-1}, and by assumption, a maximum legal matching of G'_{i-1}. By Lemma 3, the edges we delete when constructing G'_i belong to no maximum matching of G'_{i-1}, in particular $M \cap E_{\leq i-1}$. So $M \cap E_{\leq i-1}$ is contained in G'_i. Furthermore, since M is a matching and $M \supseteq M \cap E_{\leq i-1}$, it follows that M contains no deleted edges, and therefore must be contained in G'_i. □

Lemma 6. *Let M_i and M_j be the matchings produced by the algorithm, where $i < j$. Then M_i and M_j have the same number of edges with rank at most i.*

Proof. M_i consists of edges contained within top-level supervertices of G'_i, and edges between top-level (super)vertices of G'_i. We have already shown that augmenting through a supervertex does not change the number of matching edges internal to the supervertex. Hence, M_j contains the same number of such edges as M_i.

By Lemma 3, the remaining edges of M_i are all be incident on some GED-O$[G'_i]$ or GED-P$[G'_i]$ (super)vertex. Since these vertices are matched in M_i, they are also matched in M_j, as augmenting does not affect the matched status of a vertex. Also, no edges of rank worse than i are incident on such vertices, due to deletions, and so each must be matched along a rank-i edge or better in M_j. Hence $|M_i| \leq |M_j \cap E_{\leq i}|$. Of course, $|M_j \cap E_{\leq i}| \leq |M_i|$, since all edges from $E_{\leq i}$ in G'_j are also in G'_i, and M_i is a maximum legal matching of G'_i. □

Lemma 7. *For every i, the following statements hold: 1) Every rank-maximal matching of G_i is a maximum legal matching of G'_i, and 2) M_i is a rank-maximal matching of G_i.*

Proof. For the base case, rank-maximal matchings are maximum matchings on rank-1 edges, and so both statements hold for $i = 1$. Now, by Lemma 5 and the inductive hypothesis, every rank-maximal matching of G_i is contained in G'_i. Let $\langle s_1, s_2, .., s_i \rangle$ be the signature of such a matching. By Lemma 6, M_i has the same signature as M_{i-1} up to rank-$(i-1)$. Hence, M_i's signature is $\langle s_1, s_2, .., t_i \rangle$ for some $t_i \leq s_i$, since M_{i-1} is a rank-maximal matching of G_{i-1}. However, M_i is a maximum legal matching of G'_i, hence $t_i = s_i$ and M_i is rank-maximal matching of G_i. This proves the second statement.

Now, for the first statement, let N_i be any rank-maximal matching of G_i. By Lemma 5 and the inductive hypothesis, we know that N_i is contained in G'_i. N_i has signature $\langle s_1, s_2, ..., s_i \rangle$, which is the same signature as M_i. Hence, N_i is also a maximum legal matching of G'_i. □

We now comment on the runtime of the algorithm. In each iteration i, it is clear that computing the decomposition (given a maximum matching), deleting edges and shrinking components all take $O(m)$ time. Constructing M_i from M_{i-1} requires $|M_i| - |M_{i-1}| + 1$ augmentations. At the top-level of augmenting (when supervertices are regarded as vertices), we can use the Micali and Vazirani non-bipartite matching algorithm, which runs in time $O(min(\sqrt{n}, |M_{i+1}| - |M_i| + 1)m)$. Next, we have to expand each augmenting path P through its incident supervertices. Let u be the first vertex of some supervertex C that P enters along an unmatched edge. It is clear that we can do this expansion in time linear in the size of C by appending a dummy unmatched vertex d to u, and then looking for an augmenting path from d to r in C. Since each supervertex belongs to at most one augmenting path in each round of the Micali and Vazirani algorithm, this does not affect the asymptotic runtime. It follows that after R iterations, the running time is at most $O(\min(n+R, R\sqrt{n})m)$. Using the idea in [18], we can stop once R is the rank of the worst-ranked edge in a rank-maximal matching, because we can test in $O(m)$ time if M_R is a maximum matching of G_R together with all undeleted edges of rank worse than R (in which case M_R is rank-maximal).

Theorem 1. *Let R be the rank of the worst-ranked edge in a rank-maximal matching of $G = (V, E_1 \cup \ldots \cup E_m)$. Then a rank-maximal matching of G can be found in time $O(\min(n+R, R\sqrt{n})m)$.*

3 Hardness Results

In this section we establish several NP-hardness results for a special case of SR-GRP. We refer to this restriction as STABLE MARRIAGE WITH SYMMETRIC PREFERENCES (SM-SYM). An instance of SM-SYM is an instance of SR in which the underlying graph is bipartite (with *men* and *women* representing the two sets of agents in the bipartition) subject to the restriction that a woman w_j appears in the kth tie in a man m_i's list if and only if m_i appears in the kth tie in w_j's list. Clearly an instance of SM-SYM is a bipartite instance of SR-GRP in which $rank(\{m_i, w_j\}) = k$ if and only if w_j appears in the kth tie in m_i's preference list, for any man m_i and woman w_j. Indeed it will be helpful to assume subsequently that *rank* is defined implicitly in this way, given an instance of SM-SYM.

Our first result demonstrates the NP-completeness of COM-SM-SYM, which is the problem of deciding whether a complete weakly stable matching (i.e. a weakly stable matching in which everyone is matched) exists, given an instance of SM-SYM. Our transformation begins from EXACT-MM, which is the problem of deciding, given a graph G and an integer K, whether G admits a maximal matching of size K.

Theorem 2. COM-SM-SYM *is NP-complete.*

Proof. Clearly COM-SM-SYM is in NP. To show NP-hardness, we reduce from EXACT-MM in subdivision graphs, which is NP-complete [21]. Let $G = (V, E)$, a subdivision graph of some graph G', and K, a positive integer, be an instance of EXACT-MM. Suppose that $V = U \cup W$ is a bipartition of G, where $U = \{m_1, m_2, \ldots, m_{n_1}\}$ and $W = \{w_1, w_2, \ldots, w_{n_2}\}$. Then we denote the set of

Men's preferences

$m_i : (W_i)\ (y_1\ \ y_2\ \ \ldots\ \ y_{n_1-K})$		$(1 \leq i \leq n_1)$
$x_i :\ a'_i\ \ (W)$		$(1 \leq i \leq n_2 - K)$
$a_i :\ (y_i\ \ b'_i)$		$(1 \leq i \leq K)$
$b_i :\ a'_i$		$(1 \leq i \leq K)$

Women's preferences

$w_j : (U_j)\ (x_1\ \ x_2\ \ \ldots\ \ x_{n_2-K})$		$(1 \leq j \leq n_2)$
$y_j :\ a_j\ \ (U)$		$(1 \leq j \leq n_1 - K)$
$a'_j :\ (x_j\ \ b_j)$		$(1 \leq j \leq K)$
$b'_j :\ a_j$		$(1 \leq j \leq K)$

Fig. 3. Preference lists for the constructed instance of COM-SM-SYM

vertices adjacent to a vertex $m_i \in U$ in G by W_i and similarly the set of vertices adjacent to $w_j \in W$ in G by U_j.

We construct an instance I of COM-SM-SYM as follows: let $U \cup X \cup A \cup B$ be the set of men and $W \cup Y \cup A' \cup B'$ be the set of women, where $X = \{x_1, x_2, \ldots, x_{n_2-K}\}$, $Y = \{y_1, y_2, \ldots, y_{n_1-K}\}$, $A = \{a_1, a_2, \ldots, a_K\}$, $B = \{b_1, b_2, \ldots, b_K\}$, $A' = \{a'_1, a'_2, \ldots, a'_K\}$ and $B' = \{b'_1, b'_2, \ldots, b'_K\}$. The preference lists of I are shown in Figure 3 (entries in round brackets are tied). It may be verified that I is an instance of SM-SYM. We claim that G has an exact maximal matching of size K if and only if I admits a complete weakly stable matching.

Suppose G has a maximal matching M, where $|M| = K$. We construct a matching M' in I as follows. Initially let $M' = M$. There remain $n_1 - K$ men in U that are not assigned to women in W; denote these men by m_{k_i} $(1 \leq i \leq n_1 - K)$ and add (m_{k_i}, y_i) to M'. Similarly there remain $n_2 - K$ women in W that are not assigned to men in U; denote these women by w_{l_j} $(1 \leq j \leq n_2 - K)$, and add (x_j, w_{l_j}) to M'. Finally we add (a_i, b'_i) and (b_i, a'_i) $(1 \leq i \leq K)$ to M'. It may then be verified that M' is a complete weakly stable matching in I.

Conversely suppose that M' is a complete weakly stable matching in I. Let $M = M' \cap E$. We now show that $|M| = K$. First suppose that $|M| < K$. Then since M' is a complete weakly stable matching, at least $n_1 - K + 1$ men in U must be assigned in M' to women in Y, which is impossible as there are only $n_1 - K$ women in Y. Now suppose $|M| > K$. Hence at most $n_1 - K - 1$ women in Y are assigned in M' to men in U. Then since M' is complete, there exists at least one women in Y assigned in M' to a man in A. Thus at most $K - 1$ men in A are assigned in M' to women in B'. Hence only $K - 1$ women in B' are assigned in M', contradicting the fact that M' is a compete weakly stable matching. Finally, it is straightforward to verify that M is maximal in G. ◻

The following corollary (see [3] for the proof) will be useful for establishing subsequent results.

Corollary 1. COM-SM-SYM *is NP-complete, even if each preference list comprises exactly two ties (where a tie can be of length 1).*

We next consider *minimum regret* and *egalitarian* weakly stable matchings, given an instance I of SMC-SYM, which is the restriction of SM-SYM in which each person finds all members of the opposite sex acceptable. Let U and W be the set of men and women in I respectively, let M be a weakly stable matching in I, and let p be some agent in I. Then we define the *cost* of p with respect to M, denoted by $cost_M(p)$, to be $rank(p, M(p))$. Furthermore we define the *regret* of M, denoted by $r(M)$ to be $\max_{p \in U \cup W} cost_M(p)$. M has *minimum regret* if $r(M)$ is minimised over all weakly stable matchings in I. Similarly we define the *cost* of M, denoted by $c(M)$, to be $\sum_{p \in U \cup W} cost_M(p)$. M is *egalitarian* if $c(M)$ is minimised over all weakly stable matchings in I.

We define REGRET-SMC-SYM (respectively EGAL-SMC-SYM) to be the problem of deciding, given an instance I of SMC-SYM and a positive integer K, whether I admits a weakly stable matching such that $r(M) \leq K$ (respectively $c(M) \leq K$). We now show that REGRET-SMC-SYM is NP-complete.

Theorem 3. REGRET-SMC-SYM *is NP-complete.*

Proof. Clearly the problem belongs to NP. To show NP-hardness, we reduce from the restriction of COM-SM-SYM in which each person's list has exactly two ties, which is NP-complete by Corollary 1. Let I be such an instance of this problem. We form an instance I' of REGRET-SMC-SYM as follows. Initially the people and preference lists in I and I' are identical. Next, in I', each person adds a third tie in their preference list containing all members of the opposite sex that are not already contained in their first two ties. It is not difficult to verify that I admits a complete weakly stable matching if and only if I' admits a weakly stable matching M such that $r(M) \leq 2$. □

We next prove that EGAL-SMC-SYM is NP-complete, using a result of Gergely [13], shown in Theorem 4, relating to *diagonalized* Latin squares. A *transversal* of an order-n Latin square is a set S of n distinct-valued cells, no two of which are in the same row or column. A Latin square is said to be diagonalized if the main diagonal is a transversal.

Theorem 4 (Gergely [13]). *For any integer $n \geq 3$, there exists a diagonalized Latin square of order n having a transversal which has no common entry with the main diagonal.*

Theorem 5. EGAL-SMC-SYM *is NP-complete.*

Proof. Clearly EGAL-SMC-SYM is in NP. To show NP-hardness, we reduce from the restriction of COM-SM-SYM in which each person's list has exactly two ties, which is NP-complete by Corollary 1. Let I be such an instance of this problem, where $U = \{m_1, m_2, \ldots, m_n\}$ is the set of men and $W = \{w_1, w_2, \ldots, w_n\}$ is the set of women. For each man $m_i \in U$ $(1 \leq i \leq n)$ we denote the women in the first and second ties on m_i's preference list in I by $W_{i,1}$ and $W_{i,2}$ respectively, and let $W_i = W_{i,1} \cup W_{i,2}$. Similarly for each woman $w_j \in W$ $(1 \leq j \leq n)$ we denote the men in the first and second ties on w_j's preference list in I by $U_{j,1}$ and $U_{j,2}$ respectively, and let $U_j = U_{j,1} \cup U_{j,2}$.

Men's preferences
$m_i : W_{i,1} \;\; W_{i,2} \;\; (y_1 \;\; q) \;\; y_2 \;\; \dots \;\; y_n \;\; (W \setminus W_i) \qquad (1 \le i \le n)$
$x_1 : \; y_1 \;\; q \;\; (W) \;\; y_{s_{1,2}} \;\; y_{s_{1,3}} \;\; y_{s_{1,4}} \;\; \cdots \;\; y_{s_{1,n}}$
$x_2 : \; y_2 \;\; q \;\; y_{s_{2,1}} \;\; (W) \;\; y_{s_{2,3}} \;\; y_{s_{2,4}} \;\; \cdots \;\; y_{s_{2,n}}$
$x_3 : \; y_3 \;\; q \;\; y_{s_{3,1}} \;\; y_{s_{3,2}} \;\; (W) \;\; y_{s_{3,4}} \;\; \cdots \;\; y_{s_{3,n}}$
$\vdots$
$x_n : \; y_n \;\; q \;\; y_{s_{n,1}} \;\; y_{s_{n,2}} \;\; y_{s_{n,3}} \;\; y_{s_{n,4}} \;\; \cdots \;\; (W)$
$p : \;\; q \;\; (Y) \;\; (W)$

Women's preferences
$w_j : U_{j,1} \;\; U_{j,2} \;\; (x_1 \;\; p) \;\; x_2 \;\; \dots \;\; x_n \;\; (U \setminus U_j) \qquad (1 \le j \le n)$
$y_1 : \; x_1 \;\; p \;\; (U) \;\; x_{t_{1,2}} \;\; x_{t_{1,3}} \;\; x_{t_{1,4}} \;\; \cdots \;\; x_{t_{1,n}}$
$y_2 : \; x_2 \;\; p \;\; x_{t_{2,1}} \;\; (U) \;\; x_{t_{2,3}} \;\; x_{t_{2,4}} \;\; \cdots \;\; x_{t_{2,n}}$
$y_3 : \; x_3 \;\; p \;\; x_{t_{3,1}} \;\; x_{t_{3,2}} \;\; (U) \;\; x_{t_{3,4}} \;\; \cdots \;\; x_{t_{3,n}}$
$\vdots$
$y_n : \; x_n \;\; p \;\; x_{t_{n,1}} \;\; x_{t_{n,2}} \;\; x_{t_{n,3}} \;\; x_{t_{n,4}} \;\; \cdots \;\; (U)$
$q : \;\; p \;\; (X) \;\; (U)$

Fig. 4. Preference lists for the constructed instance of EGAL-SMC-SYM

We construct an instance I' of EGAL-SMC-SYM as follows: let $U \cup X \cup \{p\}$ be the set of men and let $W \cup Y \cup \{q\}$ be the set of women, where $X = \{x_1, x_2, \ldots, x_n\}$ and $Y = \{y_1, y_2, \ldots, y_n\}$. Then we construct the preference lists in I' by considering the diagonalized Latin square $S = (s_{i,j})$ of order n, as constructed using Gergely's method [13] (we note that Gergely's method is polynomial-time computable). Without loss of generality we may assume that the entries in the main diagonal are in the order $1, 2, \ldots, n$ (this can be achieved by simply permuting symbols in S if necessary). Next we construct a matrix $T = (t_{i,j})$ from S as follows: for each i and j ($1 \le i, j \le n$), if $s_{i,j} = k$ then $t_{k,j} = i$. We claim that T is a Latin square.

For, suppose not. First suppose $t_{i,j} = t_{i,k} = l$, for some $j \neq k$. Thus it follows that $s_{l,j} = s_{l,k} = i$, contradicting the fact that S is a Latin square. Now suppose $t_{i,j} = t_{k,j} = l$, for some $i \neq k$. Therefore $s_{l,j} = i$ and $s_{l,j} = k$, which is impossible. Hence T is a Latin square. Moreover the elements $1, 2, \ldots, n$ appear in order on the main diagonal of T.

We then use S and T to constructed the preference lists as shown in Figure 4. By the construction of T from S and by inspection of the remaining preference list entries, we observe that I' is an instance of EGAL-SMC-SYM. Let $K = 2(3n + 1)$. It may be verified (see [3] for the proof) that I has a complete weakly stable matching M if and only if I' has a weakly stable matching M' such that $c(M') \le K$. □

Our final hardness result (whose proof appears in [3]) applies to SM-GRP, which is the restriction of SR-GRP to bipartite graphs. Recall that a strongly stable matching has no weakly blocking pairs. MIN-BP-SM-GRP is the problem of finding a weakly stable matching (which by definition has no strongly blocking pairs) with the minimum number of weakly blocking pairs, given an instance of SM-GRP.

Theorem 6. MIN-BP-SM-GRP *is not approximate within a factor of* $n^{1-\varepsilon}$*, for any* $\varepsilon > 0$*, unless P=NP, where* n *is the number of men and women.*

4 Future Work

We conclude with an open problem. A matching M' is *more popular than* another M if more agents prefer M' to M than M to M'. A matching M is *popular* if there is no matching M' that is more popular than it. Because the *more popular than* relation is not acyclic, popular matchings may not exist. As with rank-maximality, the problem of finding popular matchings (or proving no such matching exists) has been solved in the bipartite setting [2]. This setting involves allocating items to agents, when only agents have preferences. However, the original popular matching problem, as proposed by Gärdenfors [12], applied to the stable marriage setting (with preferences on both sides). We believe that SR-GRP, and its bipartite restriction, are promising models in which to begin to solve Gärdenfors' original problem.

Acknowledgement

We would like to thank Péter Biró and Utku Ünver for helpful remarks concerning relationships between SR-GRP and SR-GAP.

References

1. D.J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 295–304, 2007.
2. D.J. Abraham, R.W. Irving, K. Telikepalli, and K. Mehlhorn. Popular matchings. In *Proceedings of SODA '05: the 16th ACM-SIAM Symposium on Discrete Algorithms*, pages 424–432. ACM-SIAM, 2005.
3. D.J. Abraham, A. Levavi, D.F. Manlove, and G. O'Malley. The stable roommates problem with globally-ranked pairs. Technical Report TR-2007-257, University of Glasgow, Department of Computing Science, September 2007.
4. E.M. Arkin, A. Efrat, J.S.B. Mitchell, and V. Polishchuk. Geometric Stable Roommates. Manuscript, 2007. `http://www.ams.sunysb.edu/~kotya/pages/geomSR.pdf`.
5. J.J. Bartholdi and M.A. Trick. Stable matchings with preferences derived from a psychological model. *Operations Research Letters*, 5:165–169, 1986.
6. K.S. Chung. On the existence of stable roommate matchings. *Games and Economic Behavior*, 33(2):206–230, 2000.

7. F.L. Delmonico. Exchanging kidneys - advances in living-donor transplantation. *New England Journal of Medicine*, 350:1812–1814, 2004.
8. J. Edmonds. Path, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
9. T. Feder. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences*, 45:233–284, 1992.
10. H.N. Gabow and R.E. Tarjan. Faster scaling algorithms for general graph matching problems. *Journal of the ACM*, 38(4):815–853, 1991.
11. D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
12. P. Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioural Sciences*, 20:166–173, 1975.
13. E. Gergely. A simple method for constructing doubly diagonalized latin squares. *Journal of Combinatorial Theory, Series A*, 16(2):266–272, 1974.
14. D.W. Gjertson and J.M. Cecka. Living unrelated donor kidney transplantation. *Kidney International*, 58:491–499, 2000.
15. D. Gusfield and R.W. Irving. *The Stable Marriage Problem: Structure and Algorithms.* MIT Press, 1989.
16. R.W. Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6:577–595, 1985.
17. R.W. Irving and D.F. Manlove. The Stable Roommates Problem with Ties. *Journal of Algorithms*, 43:85–105, 2002.
18. R.W. Irving, D. Michail, K. Mehlhorn, K. Paluch, and K. Telikepalli. Rank-maximal matchings. *ACM Transactions on Algorithms*, 2(4):602–610, 2006.
19. L. Lovász and M.D. Plummer. *Matching Theory.* Number 29 in Annals of Discrete Mathematics. North-Holland, 1986.
20. D.F. Manlove, R.W. Irving, K. Iwama, S. Miyazaki, and Y. Morita. Hard variants of stable marriage. *Theoretical Computer Science*, 276(1-2):261–279, 2002.
21. D.F. Manlove and G. O'Malley. Student project allocation with preferences over projects. In *Proceedings of ACiD 2005: the 1st Algorithms and Complexity in Durham workshop*, volume 4 of *Texts in Algorithmics*, pages 69–80. KCL Publications, 2005.
22. K. Mehlhorn and D. Michail. Network problems with non-polynomial weights and applications. *Unpublished manuscript.*
23. 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.
24. E. Ronn. NP-complete stable matching problems. *Journal of Algorithms*, 11:285–304, 1990.
25. A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151–188, 2005.
26. A.E. Roth, T. Sönmez, and M.U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.
27. S. Scott. *A study of stable marriage problems with ties.* PhD thesis, University of Glasgow, Department of Computing Science, 2005.
28. J.J.M. Tan. A necessary and sufficient condition for the existence of a complete stable matching. *J. Algorithms*, 12(1):154–178, 1991.
29. K. Telikepalli and C. Shah. Efficient algorithms for weighted rank-maximal matchings and related problems. In *Proceedings of ISAAC '06: the 17th International Symposium on Algorithms and Computation*, pages 153–162, 2006.