Abraham, D.J. and Biro, P. and Manlove, D.F. (2006) "Almost stable" matchings in the Roommates problem. In, Fleischer, R. and Trippen, G., Eds. *Proceedings of WAOA 2005: the 3rd International Workshop on Approximation and Online Algorithms, 6-7 October, 2005* Lecture Notes in Computer Science Vol 3879, pages pp. 1-14, Palma, Mallorca.

"Almost Stable" Matchings in the Roommates Problem

David J. Abraham[1,⋆], Péter Biró[2,⋆⋆] and David F. Manlove[3,⋆⋆⋆]

[1] Computer Science Department, Carnegie-Mellon University, USA.
[2] Department of Algebra, and Department of Computer Science and Information Theory, Budapest University of Technology and Economics, Hungary.
[3] Department of Computing Science, University of Glasgow, UK.

Abstract. An instance of the classical Stable Roommates problem (SR) need not admit a stable matching. This motivates the problem of finding a matching that is "as stable as possible", i.e. admits the fewest number of blocking pairs. In this paper we prove that, given an SR instance with n agents, in which all preference lists are complete, the problem of finding a matching with the fewest number of blocking pairs is NP-hard and not approximable within $n^{\frac{1}{2}-\varepsilon}$, for any $\varepsilon > 0$, unless P=NP. If the preference lists contain ties, we improve this result to $n^{1-\varepsilon}$. Also, we show that, given an integer K and an SR instance I in which all preference lists are complete, the problem of deciding whether I admits a matching with exactly K blocking pairs is NP-complete. By contrast, if K is constant, we give a polynomial-time algorithm that finds a matching with at most (or exactly) K blocking pairs, or reports that no such matching exists. Finally, we give upper and lower bounds for the minimum number of blocking pairs over all matchings in terms of some properties of a stable partition, given an SR instance I.

1 Introduction

The Stable Roommates problem (SR) is a classical combinatorial problem that has been studied extensively in the literature [3, 9, 7, 4, 15, 8]. An instance I of SR contains an undirected 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*. (Note that the

⋆ Part of this work was done whilst at Department of Computing Science, University of Glasgow, and Max-Planck-Institut für Informatik. Email: `dabraham@cs.cmu.edu`.

⋆⋆ Partially supported by the Center for Applied Mathematics and Computational Physics, and by the Hungarian National Science Fund (grant OTKA F 037301). Email: `pbiro@cs.bme.hu`.

⋆⋆⋆ Supported by the Engineering and Physical Sciences Research Council (grant GR/R84597/01), and by Royal Society of Edinburgh/Scottish Executive Personal Research Fellowship. Email: `davidm@dcs.gla.ac.uk`.

acceptability relation is symmetric, i.e. $a_j \in A_i$ if and only if $a_i \in A_j$.) Moreover we assume that in I, a_i has a linear order $\prec_{a_i}$ over A_i, which we refer to as a_i's *preference list*. If $a_j \prec_{a_i} a_k$, we say that a_i *prefers* a_j to a_k. Given $a_j \in A_i$, define $rank_{a_i}(a_j) = 1 + |\{a_k \in A_i : a_k \prec_{a_i} a_j\}|$.

Let M be a matching in I. If $\{a_i, a_j\} \in M$, 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 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 $M(a_j)$. Let $bp_I(M)$ denote the set of blocking pairs with respect to M in I (we omit the subscript if the instance is clear from the context). Matching M is *stable* in I if $bp_I(M) = \emptyset$.

Gale and Shapley [3] showed that an instance of SR need not admit a stable matching (see for example the SR instance I_r in Figure 1 where $r = 1$). Irving [7] gave an $O(m)$ algorithm that finds a stable matching or reports that none exists, given an instance I of SR. The algorithm in [7] 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 generalise the algorithm to the problem model defined here (i.e. the case of *incomplete lists*) [4]. Henceforth we denote by SRC the special case of SR in which all preference lists are complete.

As the problem name suggests, an application of SR arises in the context of campus accommodation allocation, where we seek to assign students to share two-person rooms, based on their preferences over one another. Another application occurs in the context of forming pairings of players for chess tournaments [10]. Very recently, a more serious application of SR has been studied, involving pairwise kidney exchange between incompatible patient-donor pairs [14]. Here, preference lists can be constructed on the basis of compatibility profiles between patients and potential donors.

Empirical results [12] suggest that, as n increases, the probability that a random SR instance with n agents admits a stable matching decreases steeply. Equivalently, as n grows large, these results suggest that an arbitrary matching in a random SR instance with n agents is likely to admit at least one blocking pair. In practical situations, a blocking pair $\{a_i, a_j\}$ of a given matching M need not always lead to M being undermined by a_i and a_j, since these agents might not realise that together they block M. For example, in situations where preference lists are not public knowledge, there may be limited channels of communication that would lead to the awareness of blocking pairs in practice. Nevertheless, it is reasonable to assert that the greater the number of blocking pairs of a given matching M, the greater the likelihood that M would be undermined by a pair of agents in practice. Hence, given an SR instance that does not admit a stable matching, one may regard a matching that admits 1 blocking pair as being "more stable" than a matching that admits 10 blocking pairs, for example. This motivates the problem of finding, given an SR instance I with no stable matching, a matching in I that admits the fewest number of blocking pairs [11, 2]. Such a matching is, in the sense described here, "as stable as possible".

Given an SR instance I, define $bp(I) = \min\{|bp_I(M)| : M \text{ is a matching in } I\}$. Define MIN-BP-SR to be the problem of finding, given an SR instance I, a matching M in I such that $bp(M) = bp(I)$. (Note that, if I is an SRC instance where n is even, clearly M must be a perfect matching in I.) In Section 2, we show that MIN-BP-SR is NP-hard and very difficult to approximate. In particular we show that MIN-BP-SR is not approximable within $n^{\frac{1}{2}-\varepsilon}$, for any $\varepsilon > 0$, unless P=NP. The result holds even for complete preference lists.

We also consider the variant SRT of SR in which preference lists may include ties. Ties arise naturally in practical applications: for example in the kidney exchange context, two donors may be equally compatible for a given patient. We also denote by SRTC the special case of SRT in which all preference lists are complete. The definition of a blocking pair in the SRT and SRTC cases is identical to that given for SR (however the term "prefers" in the SR definition is interpreted as "strictly prefers" in the presence of ties). (Note that in [8], stable matchings in SRT and SRTC are referred to as *weakly stable* matchings, where three stability definitions are given; however weak stability is the more commonly-studied notion in the literature.) Clearly an instance of SRTC need not admit a stable matching. Moreover it is known [13, 8] that the problem of deciding whether a stable matching exists, given an instance of SRTC, is NP-complete. Let MIN-BP-SRT denote the variant of MIN-BP-SR in which preference lists may include ties. In Section 2, we show that MIN-BP-SRT is not approximable within $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless P=NP. The result holds even if all preference lists are complete, there is at most one tie per list, and each tie has length 2.

We now remark on the format of the inapproximability results that we present for MIN-BP-SR and MIN-BP-SRT. We implicitly assume that a given instance I of the former problem is unsolvable, so that $bp(I) \geq 1$. Recall that the solvability or otherwise of I can be determined in $O(m)$ time [7, 4]. Hence $bp(I)$ can be regarded as the objective function for measuring performance guarantee. On the other hand, given an instance I of MIN-BP-SRT, we do not assume that I is unsolvable, since the problem of deciding whether this is the case is NP-complete [13, 8]. Hence possibly $bp(I) = 0$, and therefore we use $opt(I)$ to measure performance guarantee, where $opt(I) = 1 + bp(I)$. In fact our inapproximability result for MIN-BP-SRT shows that, given any $\varepsilon > 0$, it is NP-hard to distinguish between the cases that I admits a stable matching, and $bp(I) \geq n^{1-\varepsilon}$.

We also consider the case that we require a matching to admit *exactly* K blocking pairs. Define EXACT-BP-SR to be the problem of deciding, given an SR instance I and an integer K, whether I admits a matching M such that $bp(M) = K$. In Section 2 we show that EXACT-BP-SR is NP-complete (even for complete preference lists). However by contrast, in Section 3, we prove that EXACT-BP-SR is solvable in polynomial time if K is a constant. In particular we give an $O(m^{K+1})$ algorithm that takes as input an SR instance I and a constant integer K, and finds a matching M in I such that $bp(M) = K$, or reports that no such matching exists. We show how to adapt this algorithm to find a matching M in I such that $bp(M) \leq K$, or report that no such matching exists.

We next give a remark regarding related work. An alternative method has been considered in the literature for coping with instances of SR that do not admit a stable matching. Tan [16] defined a *stable partition* in a given instance I of SR, which is a generalisation of the concept of a stable matching in I. Following

$a_{4i+1} : a_{4i+2} \quad a_{4i+3} \quad a_{4i+4}$ $\qquad (0 \leq i \leq r-1)$
$a_{4i+2} : a_{4i+3} \quad a_{4i+1} \quad a_{4i+4}$
$a_{4i+3} : a_{4i+1} \quad a_{4i+2} \quad a_{4i+4}$ $\qquad M_r^1 = \{\{a_{4i+1}, a_{4i+2}\} : 0 \leq i \leq r-1\}$
$a_{4i+4} : a_{4i+1} \quad a_{4i+2} \quad a_{4i+3}$ $\qquad M_r^2 = M_r^1 \cup \{\{a_{4i+3}, a_{4i+4}\} : 0 \leq i \leq r-1\}$

Fig. 1. Instance I_r of SR and two matchings M_r^1, M_r^2 in I_r.

[12], a stable partition is a permutation Π of A satisfying the following two properties (which implicitly assume that if a_i is a fixed point of Π then a_i is appended to his own preference list):

(i) for each $a_i \in A$, a_i does not prefer $\Pi^{-1}(a_i)$ to $\Pi(a_i)$;
(ii) if a_i prefers a_j to $\Pi^{-1}(a_i)$ then a_j does not prefer a_i to $\Pi^{-1}(a_j)$.

Tan [16] showed that every instance I of SR admits a stable partition, and he also gave an $O(n^2)$ algorithm for finding such a structure in I. Moreover, starting from a stable partition, Tan [17] showed how to construct, also in $O(n^2)$ time, a largest matching M in I with the property that the matched pairs in M are stable *within themselves*. However such a matching may only be half the size of a maximum (cardinality) matching in I. Yet in many applications we seek to match as many agents as possible, and as discussed above, in order to satisfy this property, in many cases a certain number of blocking pairs may be tolerated. For example, suppose that $r \geq 1$ and consider the SR instance I_r and example matchings M_r^1, M_r^2 as shown in Figure 1. Since I_r is built up from r copies of insoluble SRC instances with 4 agents, Tan's algorithm is bound to construct a matching M in I_r of size r (such as M_r^1). Any such matching M satisfies $|bp_{I_r}(M)| \geq 2r$. However M_r^2 is a solution to MIN-BP-SR in I_r, where $|M_r^2| = 2r$ and $|bp_{I_r}(M_r^2)| = r$. In particular M_r^1 is half the size of M_r^2 and admits twice as many blocking pairs.

In Section 4, for a given SR instance I, we give upper and lower bounds for $bp(I)$ in terms of some properties of a stable partition in I.

2 Inapproximability of MIN-BP-SR and MIN-BP-SRT

In this section we present reductions showing the NP-hardness and inapproximability of each of MIN-BP-SR and MIN-BP-SRT. Define 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. Our reductions utilise the NP-completeness of EXACT-MM in cubic graphs, which we now establish.

Lemma 1. 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 [6]. 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 $k < K$. We note that maximal matchings satisfy the interpolation property [5] (i.e. G has a maximal matching of size j, for $k \leq j \leq \beta(G)$) and hence G has a maximal matching of size K. The converse is clear. □

We now define some notation. Let I be an instance of SR and let A be the set of agents in I. Given $a_i \in A$, we define a set of agents $P(a_i)$ to be a *prefix* of a_i's preference list in I if $P(a_i) \subseteq A_i$ and whenever $a_j \in P(a_i)$ and a_i prefers a_k to a_j, it follows that $a_k \in P(a_i)$. The following lemma will also be required by our reduction that establishes the inapproximability of MIN-BP-SR.

Lemma 2. *Let I be an instance of* SR *with underlying graph $G = (A, E)$. Let $a_i \in A$ and let $P(a_i)$ be a prefix of a_i's preference list in I. Then, for every $k \geq 1$, there exists an instance I' of* SR *with underlying graph $G' = (A', E')$, where $A \subseteq A'$, $|A'| = |A| + 2k$ and $E \subseteq E'$, satisfying the following two properties:*

1. *if M is any matching in I in which a_i is matched and $M(a_i) \in P(a_i)$ then there is a matching M' in I' such that $M \subseteq M'$ and $bp_{I'}(M') \cap (E' \backslash E) = \emptyset$;*
2. *if M' is any matching in I' in which a_i is matched and $M'(a_i) \notin P(a_i)$, or a_i is unmatched, then $|bp_{I'}(M') \cap (E' \backslash E)| \geq k$.*

(If I is an instance of SRC *then I' is also an instance of* SRC*.)*

Proof. Let $k \geq 1$ be given. We create a set B_k of new agents, where $B_k = \{b_2, \ldots, b_{2k+1}\}$. Let $A' = A \cup B_k$. Then $|A'| = |A| + 2k$ as required. The preference list of a_i in I' is as follows:

$$a_i : [[P(a_i)]] \;\; b_2 \;\; b_3 \;\; \ldots \;\; b_{2k+1} \;\; | \;\; [[A_i \backslash P(a_i)]]$$

where, for $S \subseteq A_i$, $[[S]]$ denotes those members of S listed in the order induced from a_i's preference list in I. For exposition purposes, we also denote a_i by b_1.

For $2 \leq r \leq 2k + 1$, the preference list of b_r in I' is as follows:

$$b_r : b_{r+1} \;\; b_{r+2} \;\; \ldots \;\; b_{2k+1} \;\; b_1 \;\; b_2 \;\; \ldots \;\; b_{r-1} \;\; | \;\; \ldots$$

where $\ldots$ at the end of b_r's list denotes all agents in A in arbitrary strict order.

Let $B'_k = \{b_1\} \cup B_k$. For any agent $b_r \in B'_k$, the agents to the left of the symbol $|$ in b_r's preference list in I' are called the *proper agents* for b_r.

Finally, every agent in $A \backslash \{a_i\}$ forms a preference list in I' by appending the members of B_k to their preference list in I (in arbitrary strict order). The definition of E' follows by construction of the preference lists in I'; hence $E \subseteq E'$. Given a matching M' in I' and an agent $b_r \in B_k$ who is matched in M', define $pr(b_r, M')$ to be the set of agents whom b_r prefers to $M'(b_r)$.

To show (1) above, let M be a matching in I such that a_i is matched in M and $M(a_i) \in P(a_i)$. Let $M' = M \cup \{\{b_r, b_{k+r}\} : 2 \leq r \leq k + 1\}$. Suppose that

$\{b_r, b_s\} \in bp_{I'}(M') \cap (E' \backslash E)$, where $b_r, b_s \in B_k$ and $r < s$. We firstly suppose that $2 \le r \le k+1$. Then $M'(b_r) = b_{r+k}$. As $b_s \in pr(b_r, M') = \{b_{r+1}, \ldots, b_{r+k-1}\}$ and $|pr(b_r, M')| = k-1$, it follows that $M'(b_s) \in \{b_{r+k+1}, \ldots, b_{2k+1}, b_2, \ldots, b_{r-1}\}$, so that $b_r \notin pr(b_s, M')$, a contradiction. Now suppose that $k+2 \le r \le 2k+1$. Then $M'(b_r) = b_{r-k}$. As $b_s \in pr(b_r, M') \backslash \{b_1\} = \{b_{r+1}, \ldots, b_{2k+1}, b_2, \ldots, b_{r-k-1}\}$ and $|pr(b_r, M') \backslash \{b_1\}| = k-1$, it follows that $M'(b_s) \in \{b_{r-k+1}, \ldots, b_{r-1}\}$, so that $b_r \notin pr(b_s, M')$, a contradiction. Finally it is easy to see that $\{a_j, b_l\} \notin bp_{I'}(M') \cap (E' \backslash E)$ for any $a_j \in A$ and $b_l \in B_k$. Hence $bp_{I'}(M') \cap (E' \backslash E) = \emptyset$ as required.

To show (2) above, let M' be a matching in I', and suppose that a_i is matched in M' and $M'(a_i) \notin P(a_i)$, or a_i is unmatched in M'. Then there is an agent $b_j \in B'_k$ who is not matched to a proper agent in M'. Define E'' to be the edges in the subgraph of G' induced by B'_k. Suppose $|M' \cap E''| = t$. Then $t \le k$. Also $2(k-t)$ agents in $B'_k \backslash \{b_j\}$ are not matched to a proper agent in M'. Now suppose that $\{b_r, b_s\} \in M' \cap E''$. Then $B'_k \backslash \{b_r, b_s\} \subseteq pr(b_r, M') \cup pr(b_s, M')$. Hence either $\{b_j, b_r\}$ or $\{b_j, b_s\}$ belongs to $bp_{I'}(M') \cap (E' \backslash E)$. Now suppose that $b_r \in B'_k \backslash \{b_j\}$ is not matched to a proper agent in M'. Then $\{b_j, b_r\} \in bp_{I'}(M') \cap (E' \backslash E)$. Hence $|bp_{I'}(M') \cap (E' \backslash E)| \ge t + 2(k-t) = 2k - t \ge k$ as required. □

Henceforth we adopt the following notation, given an instance I of SR. Given an agent a_i, a prefix $P(a_i)$ of a_i's preference list and an integer $k \ge 1$, the symbol $G_k(a_i)$ in a_i's preference list following the members of $P(a_i)$ denotes the introduction of the new agents in B_k together with their preference lists, and the insertion of the members of B_k in subscript order at the relevant point in a_i's preference list, as described by the proof of Lemma 2. Given two agents a_i, a_j and integers $k, l \ge 1$, usage of the symbols $G_k(a_i)$ and $G_l(a_j)$ in the preference lists of a_i and a_j respectively implies that the agents in B_k as introduced for a_i are disjoint from the agents in B_l as introduced for a_j.

We now present a gap-introducing reduction, starting from EXACT-MM, that establishes the hardness of approximating MIN-BP-SR.

Theorem 1. MIN-BP-SR *is not approximable within* $n^{\frac{1}{2}-\varepsilon}$*, for any* $\varepsilon > 0$*, unless P=NP. The result holds even for complete preference lists.*

Proof. Let $\varepsilon > 0$ be given. Let $G = (V, E)$ (a cubic graph) and K (a positive integer) be an instance of EXACT-MM. Assume that $V = \{v_1, \ldots, v_p\}$ and $q = |E|$. We assume that $2K \le p$, for otherwise EXACT-MM trivially has a "no" answer. Let $t = \lceil \frac{1}{\varepsilon} \rceil$ and let $C = D = p^t$. For each i ($1 \le i \le p$), let $v_{j_i}, v_{k_i}, v_{l_i}$ denote the three vertices adjacent to v_i in G. For each s ($1 \le s \le 4$), let $U^s = \{u_i^s : 1 \le i \le p\}$. Let $U = \cup_{s=1}^4 U^s$, $H = \{h_1, h_2, \ldots, h_{p-2K}\}$, $X = \{x_1, x_2, \ldots, x_C\}$, $Y = \{y_1, y_2, \ldots, y_C\}$ and $Z = \{z_i^s : 1 \le i \le p \wedge 1 \le s \le 3\}$.

For each $\{v_i, v_j\} \in E$, define $\sigma_{i,j} = 1, 2, 3$ according as v_j is v_{j_i}, v_{k_i} or v_{l_i} respectively. Also define $W_{i,j}^s = \{w_{i,j}^{r,s} : 1 \le r \le C\}$ ($1 \le s \le 2$) and $W_{i,j} = W_{i,j}^1 \cup W_{i,j}^2$. (We remark that $\{v_i, v_j\}$ gives rise to both $\sigma_{i,j}$ and $\sigma_{j,i}$, and both $W_{i,j}$ and $W_{j,i}$.) Let $W = \cup_{\{v_i, v_j\} \in E} W_{i,j}$.

We create an instance I of SRC in which the set A of agents includes $U \cup Z \cup H \cup X \cup Y \cup W$ and also additional agents that arise from instances of gadgets that are constructed implicitly by the proof of Lemma 2. The preference lists of

$$
\begin{array}{ll}
u_i^1 : z_i^1 \;\; u_{j_i}^{\sigma_{j_i,i}} \;\; [W_{i,j_i}^1] \;\; [W_{i,k_i}^1] \;\; [W_{i,l_i}^1] \;\; [H] \;\; [X] \;\; \ldots & (1 \le i \le p) \\
u_i^2 : z_i^2 \;\; u_{k_i}^{\sigma_{k_i,i}} \;\; [X] \;\; \ldots & (1 \le i \le p) \\
u_i^3 : z_i^3 \;\; u_{l_i}^{\sigma_{l_i,i}} \;\; [X] \;\; \ldots & (1 \le i \le p) \\
u_i^4 : z_i^1 \;\; z_i^2 \;\; z_i^3 \;\; [X] \;\; \ldots & (1 \le i \le p) \\
z_i^s : u_i^s \;\; u_i^4 \;\; [X] \;\; \ldots & (1 \le i \le p \wedge 1 \le s \le 3) \\
h_k : [U^1] \;\; [X] \;\; \ldots & (1 \le k \le p-2K) \\
x_r : [U] \;\; [Z] \;\; [H] \;\; [W] \;\; y_r \;\; \ldots & (1 \le r \le C) \\
y_r : x_r \;\; G_D(y_r) \;\; \ldots & (1 \le r \le C) \\
w_{i,j}^{r,1} : w_{j,i}^{r,1} \;\; u_i^1 \;\; w_{j,i}^{r,2} \;\; [X] \;\; \ldots & (1 \le i < j \le p \wedge \{v_i, v_j\} \in E \wedge 1 \le r \le C) \\
w_{i,j}^{r,2} : w_{j,i}^{r,2} \;\; w_{j,i}^{r,1} \;\; [X] \;\; \ldots & (1 \le i < j \le p \wedge \{v_i, v_j\} \in E \wedge 1 \le r \le C) \\
w_{j,i}^{r,1} : w_{i,j}^{r,2} \;\; u_j^1 \;\; w_{i,j}^{r,1} \;\; [X] \;\; \ldots & (1 \le i < j \le p \wedge \{v_i, v_j\} \in E \wedge 1 \le r \le C) \\
w_{j,i}^{r,2} : w_{i,j}^{r,1} \;\; w_{i,j}^{r,2} \;\; [X] \;\; \ldots & (1 \le i < j \le p \wedge \{v_i, v_j\} \in E \wedge 1 \le r \le C)
\end{array}
$$

Fig. 2. Preference lists in the constructed SR instance I.

the agents in $U \cup Z \cup H \cup X \cup Y \cup W$ are shown in Figure 2. In a given agent a's preference list, the symbol $[S]$, for $S \subseteq U \cup Z \cup H \cup X$, denotes all members of S listed in increasing subscript order. Similarly, for $S \subseteq W$, the symbol $[S]$ denotes all members of S listed in arbitrary strict order. Also, the symbol $\ldots$ denotes all remaining agents (other than a) listed in arbitrary strict order. For certain agents in I, we now define a prefix $P(a)$ of a's preference list as follows. For each agent $a \in U \cup Z \cup H \cup W$, define $P(a)$ to be the set of agents whom a prefers to every member of X. For each agent $y_r \in Y$, define $P(y_r) = \{x_r\}$.

It may be verified that the number of agents in I is $n = 7p + p - 2K + 2C + 2CD + 4qC = 2p^{2t} + 6p^{t+1} + 2p^t + 8p - 2K$ (since G is cubic), which is polynomial in the size of the given instance of EXACT-MM.

Suppose that M is a maximal matching in G, where $|M| = K$. We create a matching M' in I as follows. Let $\{v_i, v_j\} \in E$ where $i < j$. Suppose firstly that $\{v_i, v_j\} \in M$. Let $s_1 = \sigma_{i,j}$ and let $s_2 = \sigma_{j,i}$. Add the pairs $\{u_i^{s_1}, u_j^{s_2}\}$, $\{u_i^s, z_i^s\}$ $(1 \le s \ne s_1 \le 3)$, $\{u_i^4, z_i^{s_1}\}$, $\{u_j^s, z_j^s\}$ $(1 \le s \ne s_2 \le 3)$, $\{u_j^4, z_j^{s_2}\}$, $\{w_{i,j}^{r,1}, w_{j,i}^{r,1}\}$, $\{w_{i,j}^{r,2}, w_{j,i}^{r,2}\}$ to M' $(1 \le r \le C)$. Now suppose that $\{v_i, v_j\} \notin M$. If v_j is unmatched in M, add the pairs $\{w_{i,j}^{r,1}, w_{j,i}^{r,2}\}$, $\{w_{i,j}^{r,2}, w_{j,i}^{r,1}\}$ $(1 \le r \le C)$ to M', otherwise add the pairs $\{w_{i,j}^{r,1}, w_{j,i}^{r,1}\}$, $\{w_{i,j}^{r,2}, w_{j,i}^{r,2}\}$ $(1 \le r \le C)$ to M'.

There remain $p - 2K$ agents in U^1 who are unmatched in M' – let $u_{t_1}^1, u_{t_2}^1, \ldots, u_{t_{p-2K}}^1$ denote these agents, where $t_1 < t_2 < \ldots < t_{p-2K}$. Add $\{u_{t_k}^1, h_k\}$ and $\{u_{t_k}^s, z_{t_k}^{s-1}\}$ to M' $(2 \le s \le 4,\ 1 \le k \le p - 2K)$. Next add $\{x_r, y_r\}$ to M' $(1 \le r \le C)$. Finally, since $M'(y_r) \in P(y_r)$ for each agent $y_r \in Y$, we may extend M' by adding the edges that follow from Property 1 of Lemma 2 as applied to $G_D(y_r)$.

For each i ($1 \leq i \leq p$), there exists a unique s ($1 \leq s \leq 3$) such that $\{u_i^s, z_i^s\} \in bp(M')$. It may be verified that, by the maximality of M in G, these are all the blocking pairs of M' in I, and hence $|bp(M')| = p$.

Conversely suppose that G does not admit a maximal matching of size K. Suppose for a contradiction that $bp(I) < C$. Let M' be a matching in I such that $|bp(M')| = bp(I) < C$. Clearly every agent must be matched in M', as I is an instance of SRC and n is even. Also by Property 2 of Lemma 2, it follows that $\{y_r, x_r\} \in M'$ for all $y_r \in Y$, for otherwise $|bp(M')| \geq C$, a contradiction. Hence for each $a \in U \cup Z \cup H \cup W$, it follows that $M'(a) \in P(a)$, for otherwise $\{x_r, a\} \in bp(M')$ for all $x_r \in X$, so that $|bp(M')| \geq C$, a contradiction.

Also for each i ($1 \leq i \leq p$), $\{u_i^4, z_i^{s'}\} \in M'$ for some s' ($1 \leq s' \leq 3$). It follows that $\{z_i^s, u_i^s\} \in M'$ ($1 \leq s \neq s' \leq 3$). Now suppose that $\{u_i^1, w_{i,j}^{r,1}\} \in M'$ for some i, j ($1 \leq i, j \leq p$) and r ($1 \leq r \leq C$). Then $\{w_{i,j}^{r,2}, w_{j,i}^{r,2}\} \in M'$, for otherwise $M'(w_{j,i}^{r,2}) \notin P(w_{j,i}^{r,2})$. Hence $\{w_{j,i}^{r,1}, u_j^1\} \in M'$, for otherwise $M'(w_{j,i}^{r,1}) \notin P(w_{j,i}^{r,1})$. Define

$$M = \left\{ \{v_i, v_j\} \in E : i < j \wedge \begin{pmatrix} \{u_i^{s_1}, u_j^{s_2}\} \in M' \text{ where } 1 \leq s_1, s_2 \leq 3 \vee \\ \{u_i^1, w_{i,j}^{r,1}\} \in M' \text{ where } 1 \leq r \leq C \end{pmatrix} \right\}.$$

It follows that M is a matching in G. Also each agent in H is matched in M' to an agent in U^1, so that $|M| \leq K$. But each agent $u_i^s \in U$ satisfies $M'(u_i^s) \in P(u_i^s)$, so that $|M| = K$. Now suppose that M is not maximal in G. Then there exists some edge $\{v_i, v_j\} \in E$ such that each of v_i and v_j is unmatched in M. Hence $\{u_i^1, h_k\} \in M'$ and $\{u_j^1, h_l\} \in M'$ for some $h_k, h_l \in H$. Let r ($1 \leq r \leq C$) be given. If $\{\{w_{i,j}^{r,1}, w_{j,i}^{r,1}\}, \{w_{i,j}^{r,2}, w_{j,i}^{r,2}\}\} \subseteq M'$ then $\{w_{j,i}^{r,1}, u_j^1\} \in bp(M')$. If $\{\{w_{i,j}^{r,1}, w_{j,i}^{r,2}\}, \{w_{i,j}^{r,2}, w_{j,i}^{r,1}\}\} \subseteq M'$ then $\{u_i^1, w_{i,j}^{r,1}\} \in bp(M')$. Hence $|bp(M')| \geq C$, a contradiction. Thus M is a maximal matching of size K in G, a contradiction. Hence $bp(I) \geq C = p^t$ after all.

Next we show that $p^{t-1} > n^{\frac{1}{2}-\varepsilon}$. Firstly recall that

$$n = 2p^{2t} + 6p^{t+1} + 2p^t + 8p - 2K. \tag{1}$$

As G is cubic, we may assume that $p \geq 4$. Hence Equation 1 implies that $n < 16p^{2t}$, and thus $p^{t-1} > 16^{\frac{1-t}{2t}} n^{\frac{1}{2}-\frac{1}{2t}}$. As $t \geq \frac{1}{\varepsilon}$, it follows that

$$p^{t-1} > 4^{\frac{1-t}{t}} n^{\frac{1}{2}-\frac{\varepsilon}{2}}. \tag{2}$$

But Equation 1 also implies that $n \geq p^{2t}$, since $2K \leq p$. As $p \geq 4$, it follows that $n \geq 4^{2t} \geq 4^{\frac{2(t-1)}{\varepsilon t}}$, and hence $4^{\frac{1-t}{t}} \geq n^{-\frac{\varepsilon}{2}}$. Thus by Inequality 2, it follows that $p^{t-1} > n^{\frac{1}{2}-\varepsilon}$ as required.

Hence the existence of an $(n^{\frac{1}{2}-\varepsilon})$-approximation algorithm for MIN-BP-SR implies a polynomial-time algorithm for EXACT-MM in cubic graphs. This is a contradiction to Lemma 1 unless P=NP. □

Corollary 1. EXACT-BP-SR *is NP-complete, even for complete preference lists.*

Proof. We use the same reduction as in the proof of Theorem 1 (for any $\varepsilon < 1$) and set $K' = p$. Clearly G admits a maximal matching of size K if and only if I admits a matching with exactly K' blocking pairs. □

We now consider the case where preference lists may include ties. For a given instance I of SRT, we define $opt(I) = 1 + bp(I)$ as discussed in Section 1. The following result establishes the hardness of approximating MIN-BP-SRT.

Theorem 2. MIN-BP-SRT *is not approximable within $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless P=NP. The result holds even if all preference lists are complete, there is at most one tie per list, and each tie is of length 2.*

Proof. This result follows by adapting the proof of Theorem 1; we outline only the modifications here. For the revised reduction, choose $t = \lceil \frac{2}{\varepsilon} \rceil$, $C = p$ and $D = p^t$. Let $F = p^{t-1}$. Also, for each $z_i^s \in Z$, the agents u_i^s and u_i^4 are tied in joint first place in the preference list of z_i^s. All other preference list entries are as before. We now create F copies of each agent in $a \in U \cup Z \cup H \cup W$ – each copy of a is denoted by $a(s)$ $(1 \leq s \leq F)$. In the preference list of $a(s)$ in I, we replace b by $b(s)$ for each agent $b \in U \cup Z \cup H \cup W$ who is a proper agent for a. In the preference list of each agent in X, we replace b by $b(1), \ldots, b(F)$ for each agent $b \in U \cup Z \cup H \cup W$. For each s $(1 \leq s \leq F)$, the *class of agents* $C(s)$ comprises those agents $a(s)$ such that $a \in U \cup Z \cup H \cup W$.

As in the proof of Theorem 1, if G admits a maximal matching of size K, we may construct a matching M' in I. However M' is modified as follows: if $\{a, b\} \in M'$ for $a, b \in U \cup Z \cup H \cup W$, we replace $\{a, b\}$ by $\{a(s), b(s)\}$ $(1 \leq s \leq F)$. The presence of the ties now implies that M' is stable in I, so that $opt(I) = 1$.

Conversely if G does not admit a maximal matching of size K, then as in the proof of Theorem 1, we let M' be any matching in I such that $|bp(M')| = bp(I)$. If $\{x_r, y_r\} \notin M'$ for some r $(1 \leq r \leq C)$, it follows that $|bp(M')| \geq D$. Otherwise, it may be verified that each class of agents $C(s)$ $(1 \leq s \leq F)$ contributes at least C blocking pairs of M', for if not then G admits a maximal matching of size K. Further, these F sets of blocking pairs are pairwise disjoint, so that $|bp(M')| \geq FC = D$. Hence $opt(I) \geq D + 1 = p^t + 1$.

Next we show that $p^t \geq n^{1-\varepsilon}$. For, we firstly note that $n = (8p - 2K + 4qC)F + 2C + 2CD$, so that

$$n = 8p^{t+1} + 8p^t - 2Kp^{t-1} + 2p. \qquad (3)$$

Without loss of generality we may assume that $p \geq 9$. Hence Equation 3 implies that $n \leq 9p^{t+1}$, and thus $p^t \geq 9^{-\frac{t}{t+1}} n^{1-\frac{1}{t+1}}$. As $t \geq \frac{2}{\varepsilon}$, it follows that

$$p^t \geq 9^{-\frac{t}{t+1}} n^{1-\frac{\varepsilon}{2}}. \qquad (4)$$

Equation 3 also implies that $n \geq 9^t$, since $2K \leq p$. It follows that $n \geq 9^{\frac{2t}{\varepsilon(t+1)}}$, and hence $9^{-\frac{t}{t+1}} \geq n^{-\frac{\varepsilon}{2}}$. Thus by Inequality 4, it follows that $p^t \geq n^{1-\varepsilon}$ as required.

Hence the existence of an $(n^{1-\varepsilon})$-approximation algorithm for MIN-BP-SRT implies a polynomial-time algorithm for EXACT-MM in cubic graphs. This is a contradiction to Lemma 1 unless P=NP. □

We denote by EXACT-BP-SRT the extension of EXACT-BP-SR to the SRT case. Corollary 1 may be strengthened for EXACT-BP-SRT as follows. It is known that the problem of deciding whether an SRTC instance I admits a stable matching is NP-complete [13, 8]. Form an SRTC instance J by adding to I a new agent a_i such that $A_i = A \backslash \{a_i\}$ and $P(a_i) = \emptyset$, together with the new agents that are created by Lemma 2 as applied to a_i, with $k = K$. Clearly I admits a stable matching if and only if J admits a matching with exactly K blocking pairs. We have therefore proved:

Theorem 3. EXACT-BP-SRT *is NP-complete for each fixed* $K \geq 0$.

3 Polynomial-time algorithm for fixed K

In this section we consider the case that 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 in I such that $|bp_I(M)| = K$, or reports that no such matching exists. Later, we show how to modify this algorithm if we require that $|bp_I(M)| \leq K$.

Our algorithm is based on generating subsets B of edges of G, where $|B| = K$ – 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 in M.

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 I to E_B. By construction of G_B, it is immediate that 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 [4] to I_B – suppose it finds a stable matching M in I_B. If all agents in P are matched then, as we will show, $bp_I(M) = B$, and hence $|bp_I(M)| = K$ – 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 [4] 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 3. The following theorem establishes its correctness and complexity.

```
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;
```

Fig. 3. Algorithm K-BP.

Theorem 4. *Given an* SR *instance I and a fixed constant K, Algorithm K-BP finds a matching with exactly 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 $bp_I(M) = B$. As previously mentioned, $B \subseteq bp_I(M)$. 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 $|B| = K$. 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 $|B'| = K$. 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 $e \in bp_I(M)$. As $B \cap E_B = \emptyset$, 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 [4, 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, $bp_I(M') = B$.

On the other hand suppose that there is no matching M in I such that $|bp_I(M)| = K$. 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.

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 $O(m)$. □

Note that 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) $|bp_I(M)| = K$, and (ii) M is of maximum cardinality with respect to (i). This extension uses the fact that all stable matchings in I_B have the same size [4, Theorem 4.5.2], so that the choice of stable matching constructed by the algorithm is not of significance for Condition (ii).

Finally we remark that Algorithm K-BP may easily be modified in order to find a matching M such that $bp_I(M) \leq K$: the outermost loop iterates over all subsets B of E such that $|B| \leq K$. Again, one can find a maximum such matching if required. The time complexity of the algorithm remains unchanged.

4 Upper and lower bounds for $bp(I)$

In this section we present upper and lower bounds for $bp(I)$, given an SR instance I, in terms of properties of a stable partition as defined in Section 1. The following results concerning stable partitions were established by Tan [16].

Theorem 5 ([16]). *Given an* SR *instance I,*

1. *I admits a stable partition Π, which may be found in $O(n^2)$ time;*
2. *if C_i is an odd-length cycle in Π of length ≥ 1 (henceforth an* odd cycle*) in Π then C_i is an odd cycle in any stable partition of Π;*
3. *I admits a stable matching if and only if Π has no odd cycle of length ≥ 3.*

Let $\mathcal{C}$ denote the set of odd cycles of length ≥ 3 in a stable partition Π. Given $C_i \in \mathcal{C}$, let $d_i = \min_{a_j \in C_i} d_G(a_j)$, where $d_G(a_j)$ denotes the degree of vertex a_j in the underlying graph G of I. We firstly give an upper bound for $bp(I)$.

Lemma 3. *Given an* SR *instance I, the bound $bp(I) \leq \sum_{C_i \in \mathcal{C}} (d_i - 1)$ holds.*

Proof. We firstly remark that the upper bound is invariant for I by Part 2 of Theorem 5. It follows by [17, Proposition 4.1] and [16, Proposition 3.2] that, by deleting a vertex of minimum degree from each odd cycle of $\mathcal{C}$, and then by decomposing each even length cycle into pairs, we obtain a matching M that is stable in the instance J of SR so obtained. It then follows by Properties (i) and (ii) of Π as given in Section 1 that every blocking pair of M in I involves a deleted vertex, and moreover for any deleted vertex a_i, if $\Pi(a_i) = a_j$ then $\{a_i, a_j\} \notin bp_I(M)$ since a_j prefers $M(a_j) = \Pi(a_j)$ to a_i. It follows that $|bp_I(M)| \leq \sum_{C_i \in \mathcal{C}} (d_i - 1)$. □

In order to derive our lower bound for $bp(I)$, it will be helpful to utilise a construction due to Cechlárová and Fleiner [1] which involves transforming a given SR instance I into an SR instance I_e as follows. In I_e, the preference lists of

$a_k^1 : a_k^2 \;\; a_i \;\; a_k^4$		$a_k^2 : a_k^3 \;\; a_k^1$	
$a_k^3 : a_k^6 \;\; a_k^2$		$a_k^4 : a_k^1 \;\; a_k^5$	
$a_k^5 : a_k^4 \;\; a_k^6$		$a_k^6 : a_k^5 \;\; a_j \;\; a_k^3$	

Fig. 4. Preference lists of the newly-introduced agents in I_e.

the agents in A are initially the same as the corresponding preference lists in I. We then replace each edge $e_k = \{a_i, a_j\}$ (where $i < j$) in the underlying graph of I by a 6-cycle involving vertices a_k^1, a_k^2, a_k^3, a_k^4, a_k^5, a_k^6. In a_i's preference list in I_e, a_j is replaced by a_k^1, whilst in a_j's preference list in I_e, a_i is replaced by a_k^6. The preference lists of the newly-introduced agents are shown in Figure 4.

Cechlárová and Fleiner [1] showed that a stable matching M in I corresponds to a stable matching M_e in I_e, and vice versa, as follows:

- $\{a_i, a_j\} \in M \Leftrightarrow \{a_i, a_k^1\}, \{a_k^2, a_k^3\}, \{a_k^4, a_k^5\}, \{a_k^6, a_j\} \in M_e$
- $\{a_i, a_j\} \notin M$ and a_i prefers $M(a_i)$ to $a_j \Rightarrow \{a_k^1, a_k^4\}, \{a_k^2, a_k^3\}, \{a_k^5, a_k^6\} \in M_e$
- $\{a_i, a_j\} \notin M$ and a_i prefers a_j to $M(a_i) \Rightarrow \{a_k^1, a_k^2\}, \{a_k^3, a_k^6\}, \{a_k^4, a_k^5\} \in M_e$
- $\{a_i, a_j\} \notin M \Leftarrow \{a_k^1, a_k^4\}, \{a_k^2, a_k^3\}, \{a_k^5, a_k^6\} \in M_e$ or $\{a_k^1, a_k^2\}$, $\{a_k^3, a_k^6\}$, $\{a_k^4, a_k^5\} \in M_e$

where $\{a_i, a_j\} = e_k$. Similarly, given stable partitions Π and Π_e in I and I_e respectively, we can prove that $\Pi(a_i) = a_j$ in an odd cycle if and only if, in Π_e:

- if $i < j$ then $\langle a_i, a_k^1, a_k^2, a_k^3, a_k^6, a_j \rangle$ is in an odd cycle and $\langle a_k^4, a_k^5 \rangle$ is a cycle;
- if $j < i$ then $\langle a_i, a_k^6, a_k^5, a_k^4, a_k^1, a_j \rangle$ is in an odd cycle and $\langle a_k^2, a_k^3 \rangle$ is a cycle.

Lemma 4. *Given an* SR *instance I, the bound $bp(I) \geq \left\lceil \frac{|\mathcal{C}|}{2} \right\rceil$ holds.*

Proof. It follows from the proof of Theorem 4 that $bp(I) = k$ if and only if k is the minimum number for which there exists a set S of k edges such that the SR instance I' obtained by deleting the edges in S from I admits a stable matching. To delete an edge $e_k = \{a_i, a_j\}$ from I is equivalent to deleting the two vertices a_k^1 and a_k^6 from I_e. That is, after deleting the above set S of edges, instance I' has a stable matching if and only if, after deleting the corresponding k pairs of vertices from I_e, the obtained instance I'_e has a stable matching. But by [17, Theorem 4.2], the number of odd cycles can decrease by at most one after deleting one vertex, so after deleting k edges from I, the number of odd cycles can decrease by at most $2k$ in I_e. Hence if $|\mathcal{C}| > 2k$, then I'_e still has at least one odd cycle of length ≥ 3, so neither I'_e nor I' can admit a stable matching. □

5 Concluding remarks

The strong inapproximability results presented in this paper are perhaps surprising, in view of Theorem 5 and the various structural properties of a stable partition [16, 17]. We conclude with two open problems.

Firstly, given an SR instance I and a matching M in I, it follows that $bp(M) \leq m = O(n^2)$. Is there an approximation algorithm for MIN-BP-SR with performance guarantee $o(m)$?

Secondly, it remains open to determine whether the bounds for $bp(I)$ presented in Section 4 are tight, and in particular to establish values of k_n and to obtain a characterisation of I_n such that I_n is an SR instance with n agents, in which $bp(I_n) = k_n$ and $bp(I_n)$ is maximum over all SR instances with n agents.

Acknowledgements

We would like to thank Katarína Cechlárová and Rob Irving for helpful discussions. The problem of finding a matching with at most K blocking pairs, for a fixed integer K, was suggested by Rob Irving.

References

1. K. Cechlárová and T. Fleiner. On a generalization of the stable roommates problem. *ACM Transactions on Algorithms*, 1(1):143–156, 2005.
2. K. Eriksson and P. Strimling. How unstable are matchings from decentralized mate search? Preprint, 2005. Submitted for publication.
3. D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
4. D. Gusfield and R.W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.
5. F. Harary. Maximum versus minimum invariants for graphs. *J. Graph Theory*, 7:275–284, 1983.
6. J.D. Horton and K. Kilakos. Minimum edge dominating sets. *SIAM J. Discrete Mathematics*, 6:375–387, 1993.
7. R.W. Irving. An efficient algorithm for the "stable roommates" problem. *J. Algorithms*, 6:577–595, 1985.
8. R.W. Irving and D.F. Manlove. The Stable Roommates Problem with Ties. *J. Algorithms*, 43:85–105, 2002.
9. D.E. Knuth. *Mariages Stables*, Les Presses de L'Université de Montréal, 1976.
10. E. Kujansuu, T. Lindberg, and E. Mäkinen. The stable roommates problem and chess tournament pairings. *Divulgaciones Matemáticas*, 7(1):19–28, 1999.
11. M. Niederle and A.E. Roth. Market culture: How norms governing explording offers affect market performance. NBER working paper 10256, January 2004.
12. B.G. Pittel and R.W. Irving. An upper bound for the solvability probability of a random stable roommates instance. *Rand. Struct. Algorithms*, 5(3):465–486, 1994.
13. E. Ronn. NP-complete stable matching problems. *J. Algorithms*, 11:285–304, 1990.
14. A.E. Roth, T. Sönmez, and M. Utku Ünver. Pairwise kidney exchange. To appear in *Journal of Economic Theory*.
15. A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis* Cambridge University Press, 1990.
16. J.J.M. Tan. A necessary and sufficient condition for the existence of a complete stable matching. *J. Algorithms*, 12:154–178, 1991.
17. J.J.M. Tan. Stable matchings and stable partitions. *International J. Computer Mathematics*, 39:11–20, 1991.