

Adapting Stable Matchings to Forced and Forbidden Pairs

Niclas Boehmer
Technische Universität Berlin
Berlin, Germany
niclas.boehmer@tu-berlin.de

Klaus Heeger
Technische Universität Berlin
Berlin, Germany
heeger@tu-berlin.de

ABSTRACT

We introduce the problem of adapting a stable matching to forced and forbidden pairs. Specifically, given a stable matching M_1 , a set Q of forced pairs, and a set P of forbidden pairs, we want to find a stable matching that includes all pairs from Q , no pair from P , and that is as close as possible to M_1 . We study this problem in four classical stable matching settings: STABLE ROOMMATES (WITH TIES) and STABLE MARRIAGE (WITH TIES).

As our main contribution, we employ the theory of rotations for STABLE ROOMMATES to develop a polynomial-time algorithm for adapting STABLE ROOMMATES matchings to forced pairs. In contrast to this, we show that the same problem for forbidden pairs is NP-hard. However, our polynomial-time algorithm for the case of only forced pairs can be extended to a fixed-parameter tractable algorithm with respect to the number of forbidden pairs when both forced and forbidden pairs are present. Moreover, we also study the setting where preferences contain ties. Here, depending on the chosen stability criterion, we show either that our algorithmic results can be extended or that formerly tractable problems become intractable.

KEYWORDS

Stable Marriage; Stable Roommates; forced and forbidden pairs; incremental algorithms; rotations; NP-hardness; polynomial-time algorithm; W[1]-hardness; FPT

ACM Reference Format:

Niclas Boehmer and Klaus Heeger. 2023. Adapting Stable Matchings to Forced and Forbidden Pairs. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Alice was recently hired as a tech lead and the company gave her the possibility to select her own team of software developers. After doing so, as it is a company-wide policy to use pair programming, Alice faces the problem of grouping her developers into pairs. Because Alice is a fan of stable matchings, she organizes this by asking each software developer for his or her preferences over the other developers. Subsequently, she computes and implements a stable matching (i.e., a matching where no two developers prefer each other to their assigned partner). Unfortunately, after a couple of weeks, Alice notices that Bob and Carol, who currently work together, like each other a little bit too much so that they spend most of their time not working productively. Thus, she wants to assign both of them to a different partner. In contrast, Alice learns that

Dan and Eve, who currently do not work together, have quite complementary skill sets. She believes that both of them would greatly benefit from working with each other. Now, she faces the problem of finding a new stable matching that respects her wishes. However, as she observed that most pairs initially needed some time to find a joint way of working, she wants to minimize the number of new pairs, i.e., she wants the new matching to be as close as possible to the current one.

More formally, the problem can be described as follows. Alice is given a stable matching M_1 of some agents with preferences over each other, a set Q of forced pairs (those pairs need to be included in the new matching) and a set P of forbidden pairs (none of these pairs is allowed to appear in the new matching), and she wants to find a new stable matching respecting the forced and forbidden pairs which is as close as possible to M_1 . We initiate the study of the decision variant of this problem, where we are additionally given an integer k and the symmetric difference between the old and the new matching shall be upper-bounded by k , in the following classical stable matching settings: STABLE ROOMMATES and its bipartite variant STABLE MARRIAGE, both combined with strict preferences or preferences containing ties. We refer to the resulting problems as ADAPT STABLE ROOMMATES/MARRIAGE [WITH TIES] TO FORCED AND FORBIDDEN PAIRS.¹ For all six problems arising this way, we either present a polynomial-time algorithm or prove its NP-hardness. Moreover, we provide a complete picture of the problems’ parameterized computational complexity² with respect to the problem-specific parameters $|P|$, $|Q|$, and k .

We now present some further application scenarios where matchings need to be adapted to forced or forbidden pairs. *First*, related to our initial toy example, in companies where on each day workers work in pairs, forced and forbidden pairs can be time-dependent. Assume that the company has some “default” assignment that is used on most days. On some days, however, the workers might need to do different tasks than usual and certain combinations of workers might not be able to work together because both of them lack some skill or a certain qualification, making some pairs forbidden on certain days. *Second*, forbidden pairs can also arise over time when assigning papers to reviewers, a task which might be modeled as a STABLE MARRIAGE instance. After assigning the papers to reviewers, it sometimes turns out that some reviewers have a conflict of interest (COI) with the paper they are supposed

¹We consider two notions of stability if preferences contain ties, i.e., weak and strong stability. In weak stability, an agent pair $\{a, b\}$ is blocking a matching if both strictly prefer each other to their current partner, whereas in strong stability it is sufficient if a strictly prefers b to its partner and b is indifferent between a and its partner.

²Our results here are mostly along the parameterized complexity classes FPT and W[1]. A problem is fixed-parameter tractable (in FPT) with respect to some parameter t if there is an algorithm solving every instance \mathcal{I} of the problem in $f(t) \cdot |\mathcal{I}|^{O(1)}$ time for some computable function f . Under standard complexity theoretical assumptions, problems that are W[1]-hard for some parameter do not admit an FPT algorithm with respect to this parameter.

	SM/STRONGLY SM WITH TIES	WEAKLY SM/SR WITH TIES	SR	STRONGLY SR WITH TIES
Forced	P (Pr. 1)	NP-h. and W[1]-h. wrt. k +number of ties for one pair (Pr. 2)	P (Th. 2)	P (Th. 3)
Forbidden	P (Pr. 1)		NP-h. (Th. 1)	NP-h. (Th. 1)
Forced and Forbidden	P (Pr. 1)		FPT wrt. #forbidden pairs in M_1 (Th. 2)	FPT wrt. #forbidden pairs in M_1 (Th. 3)

Table 1: Overview of our results. “ k ” denotes the allowed size of the symmetric difference between the old and new matching.

to review resulting in “forbidden” pairs. Then, a natural objective is to change as few assignments as possible to circumvent the newly discovered COIs. *Third*, minimally adapting stable matchings to forced or forbidden pairs might become necessary in centralized matching markets where some time passes between the computation of the matching and its implementation [7, 22] as during this time certain pairs can become either unavailable or forced. For instance, in New York’s school choice market, after the matching of students to schools was announced, around 10% of students decide to pay some (private) school to secure one of their seats instead of going to their assigned school Feigenbaum et al. [15] (which could be modeled as a forced pair). Moreover, students may also realize that they cannot attend their assigned school.

Related Work. Since the introduction of the STABLE MARRIAGE problem by Gale and Shapley [18], numerous facets of stable matching problems have been extensively studied in computer science and related areas (see, e.g., the monographs of Gusfield and Irving [21], Knuth [26], and Manlove [29]). Our problem combines two previously studied aspects of stable matching problems: forced and forbidden pairs, and incremental algorithms.

Dias et al. [11] initiated the study of stable matching problems with forced and forbidden pairs. The classical task here is to decide whether there is a stable matching including all forced pairs and no forbidden pair.³ While this problem can be solved in polynomial time if the preferences do not contain ties both in the roommates and marriage context [11, 16], the problem is NP-complete in the presence of ties for weak stability for marriage and roommates instances, even if there is only one forced and no forbidden pair [30] or one forbidden and no forced pair [9]. Motivated by the straightforward observation that a stable matching including all given forced pairs and no forbidden pairs might not exist, Cseh and Manlove [10] studied the problem of finding a matching minimizing the number of “violated constraints” (where a violated constraint is either a blocking pair or a forced pair not contained in the matching or a forbidden pair contained in the matching).

Our problem also has a clear “incremental” dimension in the sense that we want to make as few changes as possible to a stable matching to achieve a certain goal. Many authors have studied such incremental problems in the context of various stable matching scenarios [1, 3, 4, 6, 15, 17, 19, 31]. In the works of Bhattacharya et al. [1], Boehmer et al. [3, 4], Bredebeck et al. [6], Gajulapalli et al. [17], and Feigenbaum et al. [15], the focus lied on problems related to adapting matchings to change: We are given a (stable) matching of agents, then some type of change occurs (e.g., some agents revise their preferences or some agents get added or deleted) and a new (stable) matching shall be found. Here, as in our problem,

³Note that our problem reduces to the classical problem associated with forced and forbidden pairs if we set the allowed distance between M_1 and the matching to be found to infinity.

it is often assumed that changing a pair in the matching is costly so the new matching should be as close as possible to the old one. As a second type of incremental problems, Marx and Schlotter [31] and Gupta et al. [19] analyzed the computational complexity of problems where one is given a stable matching M and the task is to find a larger (almost) stable matching which is close to M . On a more general note, this paper fits into the stream of works on incremental combinatorial problems [1, 5, 8, 12, 23] where one aims at efficiently adapting solutions to changing inputs and requirements (in our case the requirement is that certain pairs are forbidden or forced), a core issue in modern algorithmics.

Our Contributions. We initiate the study of adapting stable matchings to forced and forbidden pairs. We consider this problem in six different settings and provide a complete dichotomy for the problems’ (parameterized) computational complexity with respect to the problem-specific parameters $|P|$, $|Q|$, and k . See Table 1 for an overview of our results.

In the first (short) part of the paper (Section 3), we consider the bipartite marriage setting. We prove that adapting to forced and forbidden pairs is polynomial-time solvable for STABLE MARRIAGES without ties and in case of ties in combination with strong stability (Proposition 1). However, in case ties in the preferences are allowed and we are searching for weakly stable matchings, we obtain NP-hardness and W[1]-hardness with respect to the summed number of ties and the allowed difference k between the old and the new matching (Proposition 2). These hardness results hold even if there is only one forced and no forbidden pair or if there is only one forbidden and no forced pair. As STABLE ROOMMATES generalizes STABLE MARRIAGE, these hardness results also hold for WEAKLY STABLE ROOMMATES WITH TIES.

In the second (main) part of the paper (Section 4), we focus on the STABLE ROOMMATES problem. Here, we first prove that in contrast to the bipartite setting, ADAPT STABLE ROOMMATES TO FORCED AND FORBIDDEN PAIRS is NP-hard, even if there are only forbidden pairs (Theorem 1). In contrast to this, the problem is fixed-parameter tractable with respect to the number of forbidden pairs (contained in the given matching; Theorem 2). In particular, if there are only forced pairs, then the problem is polynomial-time solvable. To the best of our knowledge, this is the first problem which is tractable for forced but intractable for forbidden pairs.⁴ The FPT-algorithm for adapting a STABLE ROOMMATES matching to forced and forbidden pairs is our main technical contribution. Our algorithm relies on exploiting the structure of the rotation poset for STABLE ROOMMATES instances in a clever way: For this, we observe that for most pairs there is a necessary (and a prohibited) rotation that needs to be part of (cannot be part of) a set of rotations

⁴Note that any problem involving only forced pairs can be reduced to a problem involving only forbidden pairs by setting for each forced pair $\{a, b\}$ all pairs containing a except for $\{a, b\}$ to be forbidden.

corresponding to a stable matching containing the pair. Using this, we can modify the set of rotations corresponding to the given matching to minimally change it to include all forced pairs. In fact, using some additional information, it is also possible to exclude forbidden pairs by modifying the rotation set. Note that as each forbidden pair in P requires a change in the matching M_1 , this algorithm also constitutes a fixed-parameter tractable algorithm for the allowed difference k between the old and the new matching. Lastly, we describe how our algorithm can be modified to also work for the STRONGLY SR WITH TIES problem by exploiting the more intricate structure of the rotation poset for this problem using similar ideas as for STABLE ROOMMATES (Theorem 3). Note that we do not consider the notion of super stability in the presence of ties in the following, as all our results for the problem variants without ties translate to super stability.⁵

The proofs (or their completions) of all results marked by (★) can be found in our full version [2].

2 PRELIMINARIES

In STABLE ROOMMATES (SR), we are given a set $A = \{a_1, \dots, a_{2n}\}$ of agents where each agent has a subset $\text{Ac}(a) \subseteq A \setminus \{a\}$ of agents it finds *acceptable*. We assume that acceptability is symmetric, i.e., $a \in \text{Ac}(a')$ for some $a, a' \in A$ implies that $a' \in \text{Ac}(a)$. Moreover, each agent $a \in A$ has (strict) preferences \succ_a over all agents it accepts, i.e., a total order over the agents $\text{Ac}(a)$. For agents $a, a', a'' \in A$, agent a *prefers* a' to a'' if $a' \succ_a a''$.

For a set A of agents, we use $\binom{A}{2}$ to denote the 2-element subsets of A ; abusing notation, we will call these 2-element subsets *pairs* although they are unordered. A *matching* M is a set of pairs $\{a, a'\} \in \binom{A}{2}$ with $a \in \text{Ac}(a')$ and $a' \in \text{Ac}(a)$, where each agent appears in at most one pair. An agent a is *matched* in some matching M if M contains a pair containing a . If a is not matched in M , then a is *unmatched*. A matching is *complete* if all agents are matched. For an agent $a \in A$ and a matching M , we denote by $M(a)$ the partner of a in M , i.e., $M(a) = a'$ if $\{a, a'\} \in M$. For two matchings M and M' and an agent a matched in both M and M' , we say that a *prefers* M to M' if a prefers $M(a)$ to $M'(a)$. An agent pair $\{a, a'\} \in \binom{A}{2}$ *blocks* a matching M if (i) $a \in \text{Ac}(a')$ and $a' \in \text{Ac}(a)$, (ii) a is unmatched or a prefers a' to $M(a)$, and (iii) a' is unmatched or a' prefers a to $M(a')$. A matching which is not blocked by any agent pair is called *stable*. An agent pair $\{a, a'\} \in \binom{A}{2}$ is a *stable pair* if there is a stable matching M with $\{a, a'\} \in M$. For two matchings M and M' , we denote by $M \Delta M'$ the set of pairs that only appear in one of M and M' , i.e., $M \Delta M' = \{\{a, a'\} \mid (\{a, a'\} \in M \wedge \{a, a'\} \notin M') \vee (\{a, a'\} \notin M \wedge \{a, a'\} \in M')\}$. The main problem studied in this paper is the following:

ADAPT SR TO FORCED AND FORBIDDEN PAIRS

Input: A set A of agents with strict preferences over each other, a stable matching M_1 , a set of forced pairs $Q \subseteq \binom{A}{2}$, a set of forbidden pairs $P \subseteq \binom{A}{2}$, and an integer k .

Question: Is there a stable matching M_2 with $Q \subseteq M_2$, $M_2 \cap P = \emptyset$, and $|M_1 \Delta M_2| \leq k$?

⁵ Fleiner et al. [16] give a polynomial-time reduction from SUPER STABLE ROOMMATES WITH TIES to STABLE ROOMMATES which declares some pairs that are not contained in any super stable matching to be forbidden. By adding these pairs to the given set of forbidden pairs, our algorithmic results for the case without ties can be adapted.

In SR WITH TIES, a generalization of SR, each agent $a \in A$ has weak preferences \succeq_a over all agents it accepts, i.e., \succeq_a is a weak order over the agents $\text{Ac}(a)$. For agents $a, a', a'' \in A$, agent a *weakly prefers* a' to a'' if $a' \succeq_a a''$, agent a is *indifferent* between a' and a'' (denoted as $a' \sim_a a''$) if both $a' \succeq_a a''$ and $a'' \succeq_a a'$, and a *strictly prefers* a' to a'' (denoted as $a' \succ_a a''$) if $a' \succeq_a a''$ but not $a'' \succeq_a a'$. We distinguish two different types of stability in the presence of ties: Under weak/strong stability, an agent pair $\{a, a'\} \in \binom{A}{2}$ *blocks* a matching M if (i) $a \in \text{Ac}(a')$ and $a' \in \text{Ac}(a)$, (ii) a is unmatched or a strictly prefers a' to $M(a)$ and (iii) a' is unmatched or a' strictly/weakly prefers a to $M(a')$. The problems ADAPT WEAKLY/STRONGLY SR WITH TIES TO FORCED AND FORBIDDEN PAIRS are defined analogous to ADAPT SR TO FORCED AND FORBIDDEN PAIRS, where instead of strict preferences weak preferences are given and weak, respectively, strong stability is required.

In the bipartite variant of SR called STABLE MARRIAGE (SM), the agents are partitioned into two set U and W . Following standard terminology, we call the elements from U *men* and the elements from W *women*. For each $m \in U$, we have $\text{Ac}(m) \subseteq W$ and for each $w \in W$ we have $\text{Ac}(w) \subseteq U$. Consequently, agents from one side can only be matched to and form blocking pairs with agents from the other side. All other definitions from above still apply. The ADAPT (STRONGLY/WEAKLY) SM (WITH TIES) TO FORCED AND FORBIDDEN PAIRS problems are defined analogously to the respective variants for SR (the only difference being that the given instance is “bipartite”, i.e., the set of agents can be split into two sets accepting only agents from the other set).

3 STABLE MARRIAGE

In this section, we study the problem of adapting stable matchings to forced and forbidden pairs in the bipartite marriage setting.

3.1 (Strongly) Stable Marriage

We start by analyzing the case where agents’ preferences are strict or when we are interested in strong stability in the presence of ties. We show that our problem is polynomial-time solvable in these settings by a simple reduction to the polynomial-time solvable WEIGHTED (STRONGLY) STABLE MARRIAGE (WITH TIES) problem, where we are given an SM instance and a weight function on the pairs and the task is to compute a minimum-weight stable matching:

PROPOSITION 1. ADAPT SM TO FORCED AND FORBIDDEN PAIRS and ADAPT STRONGLY SM WITH TIES TO FORCED AND FORBIDDEN PAIRS are solvable in $O(n \cdot m \log n)$ time.

PROOF. Both problems can be solved using the same approach: We assume that $P \cap Q = \emptyset$, as otherwise we have a trivial no instance. We define a weight function w as follows: For each forbidden pair $e \in P$, we set $w(e) := 3 \cdot n$. For each forced pair $e \in Q \setminus M_1$ that is not part of M_1 , we set $w(e) := 2 - 3 \cdot n$. For each forced pair $e \in Q \cap M_1$ that is part of M_1 , we set $w(e) := -3 \cdot n$. For each pair $e \in M_1 \setminus (P \cup Q)$ that is part of M_1 but neither forced nor forbidden, we set $w(e) = 0$. For each remaining pair e , we set $w(e) := 2$. We compute a minimum-weight stable matching M^* in $O(n \cdot m \log n)$ time (see [13] for strict preferences and [27] for the case of ties with strong stability). Note that $w(M^*) = 3 \cdot n \cdot (|P \cap M^*| - |M^* \cap Q|) + 2|M^* \setminus M_1| =$

$3 \cdot n(|P \cap M^*| - |M^* \cap Q|) + |M^* \Delta M_1|$ using that each stable matching has the same size by the Rural Hospitals Theorem [28, 32] (and thus $|M^*| = |M_1|$) for the second equality. Since $|M^*| \leq n$, it follows that $w(M^*) \leq -3 \cdot n \cdot |Q| + k$ if and only if $P \cap M^* = \emptyset$, $Q \subseteq M^*$, and $|M^* \Delta M_1| \leq k$. \square

3.2 Weakly Stable Marriage With Ties

In contrast to the previous polynomial-time solvability result for strict preferences and for strong stability from Section 3.1, we obtain strong intractability results if we consider weak stability. Note that for WEAKLY SM WITH TIES already deciding the existence of a stable matching containing a single forced pair [30] or a single forbidden pair [9] is NP-complete, implying that ADAPT WEAKLY SM WITH TIES TO FORCED AND FORBIDDEN PAIRS is NP-complete already if $|P| = 1$ or $|Q| = 1$. We extend this hardness by showing $W[1]$ -hardness when parameterized by the number of ties plus k .

PROPOSITION 2 (★). *ADAPT WEAKLY SM WITH TIES TO FORCED AND FORBIDDEN PAIRS restricted to instances where only agents from one side of the bipartition have ties in their preferences parameterized by the number of ties plus k is $W[1]$ -hard, even if $|Q| = 1$ and $P = \emptyset$ or $Q = \emptyset$ and $|P| = 1$.*

4 STABLE ROOMMATES

As the strong intractability results for ADAPT WEAKLY SM WITH TIES TO FORCED AND FORBIDDEN PAIRS from Section 3.2 extend to ADAPT WEAKLY SR WITH TIES TO FORCED AND FORBIDDEN PAIRS (as SR WITH TIES generalizes SM WITH TIES) in this section we focus on ADAPT (STRONGLY) SR (WITH TIES) TO FORCED AND FORBIDDEN PAIRS. We first prove in Section 4.1 that adapting a SR matching to forbidden pairs is NP-hard even without ties. Afterwards, in our core Section 4.2, we prove that ADAPT SR TO FORCED AND FORBIDDEN PAIRS parameterized by the number of forbidden pairs that appear in M_1 is fixed-parameter tractable (and thus that adapting an SR matching to forced pairs is polynomial-time solvable) by exploiting the rotation poset. Lastly, in Section 4.3, we extend this result to also work for ADAPT STRONGLY SR WITH TIES TO FORCED AND FORBIDDEN PAIRS.

4.1 NP-hardness of Adapt SR to Forbidden Pairs

In this section, we prove that in contrast to the bipartite marriage setting, ADAPT SR TO FORCED AND FORBIDDEN PAIRS (without ties) is already NP-hard (even if we only have forbidden pairs).

Theorem 1. *ADAPT SR TO FORCED AND FORBIDDEN PAIRS is NP-hard, even if $Q = \emptyset$ and $P \subseteq M_1$.*

PROOF. We reduce from the NP-hard INDEPENDENT SET problem [25]. Let (G, ℓ) be an instance of INDEPENDENT SET. For a vertex $v \in V(G)$, we denote by $N(v)$ the set of its neighbors in G . For each vertex $v \in V(G)$, the ADAPT SR TO FORCED AND FORBIDDEN PAIRS instance contains ten agents $a_1^v, \dots, a_5^v, b_1^v, \dots, b_5^v$. For each $v \in V(G)$, fix an arbitrary strict order of $\{a_2^w \mid w \in N(v)\}$ and denote this order by $[N^*(v)]$. For each $v \in V(G)$ the preferences of the respective ten agents are as follows (see also Figure 1):

$$a_1^v : b_1^v > b_2^v, \quad a_2^v : b_3^v > b_2^v > [N^*(v)] > b_1^v,$$

$$\begin{aligned} a_3^v : b_2^v > b_3^v, \quad a_4^v : b_5^v > b_3^v > b_4^v, \quad a_5^v : b_4^v > b_5^v \\ b_1^v : a_2^v > a_1^v, \quad b_2^v : a_1^v > a_2^v > a_3^v, \quad b_3^v : a_3^v > a_4^v > a_2^v, \\ b_4^v : a_4^v > a_5^v, \quad b_5^v : a_5^v > a_4^v \end{aligned}$$

Finally, we set $M_1 := \{\{a_i^v, b_i^v\} \mid i \in [5], v \in V(G)\}$, $P := \{\{a_2^v, b_2^v\} \mid v \in V(G)\}$, and $k := 8|V(G)| - 4\ell$. Note that M_1 is stable, as for each $v \in V(G)$, agents b_3^v, b_4^v, b_5^v , and a_1^v are matched to their top-choices (so they cannot be part of a blocking pair) and a_2^v and b_2^v are matched to their most preferred agents that are not listed above.

(\Rightarrow): Let X be an independent set of size ℓ in G . For a vertex $v \in V(G)$, we set $M^v := \{\{a_1^v, b_2^v\}, \{a_2^v, b_1^v\}, \{a_3^v, b_3^v\}, \{a_4^v, b_4^v\}, \{a_5^v, b_5^v\}\}$ and $\bar{M}^v := \{\{a_1^v, b_1^v\}, \{a_2^v, b_3^v\}, \{a_3^v, b_2^v\}, \{a_4^v, b_5^v\}, \{a_5^v, b_4^v\}\}$. We set $M^* := \bigcup_{v \in X} M^v \cup \bigcup_{v \in V \setminus X} \bar{M}^v$. Then $M^* \Delta M_1 = \{\{a_1^v, b_2^v\}, \{a_2^v, b_1^v\}, \{a_1^v, b_1^v\}, \{a_2^v, b_2^v\} \mid v \in X\} \cup \{\{a_2^v, b_3^v\}, \{a_3^v, b_2^v\}, \{a_4^v, b_5^v\}, \{a_5^v, b_4^v\} \mid v \in V(G) \setminus X\}$. Consequently, we have $|M^* \Delta M_1| = 4|X| + 8 \cdot (|V(G)| - |X|) = 8|V(G)| - 4\ell$. As M^* clearly does not contain any forbidden pair, it remains to show that M^* is stable.

It is straightforward to verify that no pair $\{a_i^v, b_j^v\}$ for $i, j \in [5]$ and $v \in V(G)$ is blocking. The remaining acceptable pairs are $\{a_2^v, a_2^w\}$ for some $\{v, w\} \in E(G)$. Since X is an independent set, we may assume without loss of generality that $v \notin X$. This implies that $M^*(a_2^v) = b_3^v >_{a_2^v} a_2^w$, implying that $\{a_2^v, a_2^w\}$ does not block M^* . Thus, M^* is stable.

(\Leftarrow): Let M^* be a stable matching with $|M^* \Delta M_1| \leq k = 8|V(G)| - 4\ell$ in the constructed instance. First note that the Rural Hospitals Theorem [21, Theorem 4.5.2] (which states that every stable matching matches the same set of agents) implies that every stable matching is complete in the constructed instance. Consequently, M^* does not contain a pair of the form $\{a_2^v, a_2^w\}$ (as otherwise one of b_1^v, \dots, b_5^v would be unmatched in M^*). Thus, for each $v \in V(G)$, we have $M^*(a_2^v) \in \{b_1^v, b_3^v\}$ (recall that we forbid the pair $\{a_2^v, b_2^v\}$ for all $v \in V(G)$). Note that $X := \{v \in V(G) \mid \{a_2^v, b_1^v\} \in M^*\}$ is an independent set: If $\{v, w\} \in E(G)$ for $v \neq w \in X$, then $\{a_2^v, a_2^w\}$ blocks M^* .

It remains to show that $|X| \geq \ell$. For each $v \in X$, we have $\{a_2^v, b_1^v\} \in M^*$ (by the definition of X) and $\{a_1^v, b_2^v\} \in M^*$ (as a_1^v would be unmatched otherwise). Consequently, $|(M^* \Delta M_1) \cap \{\{a_i^v, b_j^v\} : i, j \in [5]\}| \geq 4$. For each $v \in V(G) \setminus X$, by the definition of X we have $\{a_2^v, b_3^v\} \in M^*$. Moreover, note that M^* contains $\{a_1^v, b_1^v\}$ (as otherwise b_1^v would be unmatched) and $\{a_3^v, b_2^v\}$ (otherwise b_2^v would be unmatched). Further, M^* contains $\{a_4^v, b_5^v\}$ (otherwise $\{a_4^v, b_3^v\}$ would be blocking) and $\{a_5^v, b_4^v\}$ (otherwise a_5^v and b_4^v would be unmatched). Consequently, we have $|(M^* \Delta M_1) \cap \{\{a_i^v, b_j^v\} : i, j \in [5]\}| \geq 8$. Summing up, we get that $k = 8|V(G)| - 4\ell \geq |M_1 \Delta M^*| \geq 4|X| + 8(|V(G)| - |X|) = 8|V(G)| - 4|X|$, which is equivalent to $|X| \geq \ell$. \square

4.2 (FPT-)Algorithm for Adapt SR to Forced and Forbidden Pairs

In this section, we develop an FPT-algorithm for the ADAPT SR TO FORCED AND FORBIDDEN PAIRS problem parameterized by the number of forbidden pairs in M_1 (note that this algorithm is a




Figure 1: The preferences of agents a_i^v and b_i^v for some $v \in V(G)$ with different matchings highlighted in bold. $\{a_2^v, b_2^v\}$ is the forbidden pair. The preferences of the agents are encoded in the numbers on the edges: For an edge $\{a, a'\}$, the number x closer to a denotes the position in which a' appears in the preferences of a , i.e., a prefers exactly $x - 1$ agents to a' .

polynomial-time algorithm if no forbidden pairs are present). Our algorithm heavily relies on the rotation poset for STABLE ROOMMATES. Thus, we start this section by defining rotations (Section 4.2.1) and describing the high-level idea of our algorithm together with proving some useful facts concerning rotations (Section 4.2.2), before we present our algorithm (Section 4.2.3).

In the following we assume that all considered stable matchings (and in particular the initial matching M_1) are complete matchings, as we can otherwise modify the instance accordingly in $O(m)$ time.⁶

4.2.1 Rotations: Introduction. We first formally define what a rotation is, then discuss their relationship to Irving’s algorithm, and lastly identify different types of rotations.

Basic Definitions for Rotations. For an instance of SR, an *exposed rotation* is a sequence of agent pairs $(a_{i_0}, a_{j_0}), \dots, (a_{i_{r-1}}, a_{j_{r-1}})$ such that, for each $s \in [r]$, agent a_{i_s} ranks a_{j_s} first and $a_{j_{s+1}}$ second (where all indices in this paragraph are taken modulo r).⁷ *Eliminating* an exposed rotation $(a_{i_0}, a_{j_0}), \dots, (a_{i_{r-1}}, a_{j_{r-1}})$ means deleting, for all $s \in [r]$, all agents which a_{j_s} ranks after $a_{i_{s-1}}$ from the preferences of a_{j_s} . The *dual* $\bar{\varphi}$ of a rotation $\varphi = (a_{i_0}, a_{j_0}), \dots, (a_{i_{r-1}}, a_{j_{r-1}})$ is $\bar{\varphi} = (a_{j_0}, a_{i_{r-1}}), (a_{j_1}, a_{i_0}), (a_{j_2}, a_{i_1}), \dots, (a_{j_{r-1}}, a_{j_{r-2}})$. Note that the dual of the dual of a rotation is again the rotation itself.

Irving’s Algorithm. The theory of rotations is closely connected to Irving’s algorithm [24]. Irving’s algorithm constructs a stable matching in an SR instance (if it exists) in two phases. In the first phase, similar to the Gale-Shapely algorithm for SM, agents make proposals to each other, which are accepted or rejected. Doing so, certain parts of the agent’s preferences get deleted. Let P_0 be the

⁶If M_1 is not complete, let B be the set of agents unmatched in M_1 . For each agent $b \in B$, we add an agent b' to the instance which only finds b acceptable and which is added at the end of the preferences of b . Then, using the Rural Hospitals Theorem for SR [21], which states that each stable matching in a SR instance matches the same set of agents, it follows that all stable matchings in the modified instance contain pairs $\{b, b'\} \mid b \in B$. Consequently, the modified instance is equivalent to the original one.

⁷Notably, a rotation has no fixed start point, as we can start with any pair from the sequence resulting in shifted versions of the same rotation. In the following, we do not distinguish between these different shifted variants of the same rotation as they are the same for our purposes.

preference profile of the agents after the termination of Phase 1. Now Phase 2 consists of eliminating exposed rotations one after each other until no rotation is exposed anymore (note that after eliminating a rotation, some agents delete agents from their preferences, causing the set of exposed rotations to change). If no rotation is exposed, then either there is at least one agent with empty preferences, implying that no stable matching exists, or every agent has exactly one other agent left in its preferences, implying that matching the agents to the remaining agent in their preferences results in a stable matching. We call a preference profile a *stable table* if it can be derived from P_0 after successively eliminating exposed rotations. For an instance of STABLE ROOMMATES, the *rotations* are sequences of agent pairs which may arise as an exposed rotation in some execution of Irving’s algorithm (since Irving’s algorithm may eliminate any exposed rotation, different executions of Irving’s algorithm may result in different stable matchings and different exposed rotations)

(Non)-Singular Rotations and Further Definitions. Using the view of Irving’s algorithm now allows us to identify different types of rotations. A rotation φ is *nonsingular* if its dual $\bar{\varphi}$ is again a rotation. Otherwise, the rotation φ is *singular*. For two rotations φ, ρ , we say $\varphi \triangleright \rho$ (or φ *precedes* ρ) if φ must be eliminated before ρ , i.e., before we arrive at a stable table where ρ is exposed we need to eliminate φ . A set Z of rotations is *closed* if whenever $\rho \in Z$ and $\varphi \triangleright \rho$, then also $\varphi \in Z$. A set Z of rotations is *complete* if it contains all singular rotation and for each nonsingular rotation φ , it contains either φ or $\bar{\varphi}$. An agent pair is called *fixed* if it is contained in every stable matching. An agent b is a *stable partner* of agent a if there is a stable matching containing $\{a, b\}$, i.e., if $\{a, b\}$ is a stable pair.

Example 1. Consider the following instance of STABLE ROOMMATES (in fact, this is even an instance of STABLE MARRIAGE).

$$\begin{array}{ll}
 m_1 : w_1 > w_2 > w_3 & w_1 : m_2 > m_3 > m_1 \\
 m_2 : w_2 > w_3 > w_1 & w_2 : m_3 > m_1 > m_2 \\
 m_3 : w_3 > w_1 > w_2 & w_3 : m_1 > m_2 > m_3
 \end{array}$$

Phase 1 of Irving’s algorithm does not alter the preferences of this instance. Thus, the above preference profile is P_0 . In P_0 , rotations $\varphi_1 = (m_1, w_1), (m_2, w_2), (m_3, w_3)$ and $\varphi_2 = (w_1, m_2), (w_2, m_3), (w_3, m_1)$ are exposed. After eliminating φ_1 , rotations φ_2 and $\varphi_3 = (m_1, w_2), (m_2, w_3), (m_3, w_1)$ are exposed. After eliminating φ_2 , rotations φ_1 and $\varphi_4 = (w_1, m_3), (w_2, m_1), (w_3, m_2)$ are exposed. $\varphi_1, \varphi_2, \varphi_3, \varphi_4$ are the only rotations. Note that $\bar{\varphi}_1 = \varphi_4$ and $\bar{\varphi}_2 = \varphi_3$, implying that all four rotations are nonsingular and that no singular rotation exists. The rotation poset contains only the following two relations: φ_1 precedes φ_3 and φ_2 precedes φ_4 . Consequently, there are three closed and complete subsets of the rotation poset: $\{\varphi_1, \varphi_2\}$ (whose elimination results in the stable matching $\{\{m_1, w_2\}, \{m_2, w_3\}, \{m_3, w_1\}\}$), $\{\varphi_1, \varphi_3\}$ (whose elimination results in $\{\{m_1, w_3\}, \{m_2, w_1\}, \{m_3, w_2\}\}$), and $\{\varphi_2, \varphi_4\}$ (whose elimination results in $\{\{m_1, w_1\}, \{m_2, w_2\}, \{m_3, w_3\}\}$).

We continue by observing the following basic fact about rotations:

Lemma 1 ([21, p. 169 and Lemma 4.2.7]). *If rotation $\varphi = (a_{i_0}, a_{j_0}), \dots, (a_{i_{r-1}}, a_{j_{r-1}})$ is exposed in some stable table T , then a_{i_k} is the last agent in the preferences of a_{j_k} in T for each $k \in [0, r - 1]$. Eliminating φ in particular includes deleting the pair $\{a_{i_k}, a_{j_k}\}$ for each $k \in [0, r - 1]$.*

For our algorithm, we will exploit that it is possible to work on sets of rotations instead of stable matchings, as there is a bijection between closed complete subsets of rotations and stable matchings. In particular, given a closed and complete subset of rotations Z , there is an ordering of the rotations from Z such that starting with P_0 we can eliminate the exposed rotations one by one, resulting in a preference profile where the preferences of each agent a only contain the partner of a in the matching corresponding to Z (see also Example 1):

Lemma 2 ([21, Theorem 4.3.2]). *There is a bijection between closed and complete subsets of rotations and stable matchings. The bijection maps each closed and complete subset Z of rotations to the matching arising through the elimination of each rotation of Z .*

4.2.2 High-Level Idea and Useful Lemmas. The general idea behind our algorithm for ADAPT SR TO FORCED AND FORBIDDEN PAIRS is to successively alter the closed and complete set of rotations Z_1 corresponding to the given matching M_1 in order to include all forced and exclude all forbidden pairs. At the core of our algorithm lies the observation that rotations come with certain identifiable guarantees how “good” an agent is matched in a resulting stable matching: For instance, in case we eliminate an exposed rotation that makes agent c the last agent in the preferences of a (recall Lemma 1), we know that a is either matched to c or an agent it prefers to c in the corresponding stable matching. This allows one to identify, for some agent pair $\{a, b\}$ certain (prohibited) rotations that if included in a set of rotations guarantee that the pair cannot be part of the corresponding stable matching (those rotations guarantee that a is matched better than b). Conversely, there is often also a (necessary) rotation that needs to be included in a set of rotations corresponding to a stable matching containing the pair (the rotation that ensures that a is matched better than all agents to which it prefers b). These necessary and prohibited rotations then allow us to control whether pairs are (not) included in the output

stable matching. For instance, in order to ensure that all forced pairs are contained in the matching, we alter Z_1 to include all necessary and exclude all prohibited rotations of forced pairs (thereby changing Z_1 as little as possible to ensure that all forced pairs get included). For forbidden pairs, the situation will be slightly more complicated, as we can either not include the necessary rotation or include one of the prohibited rotations.

In order to identify necessary and prohibited rotations, we start by stating a useful characterization under which circumstances and agent b can become the last agent in the preferences of a in some stable table due to Gusfield [20]. For this, for an agent pair $\{a, b\}$, let $\rho^{a,b}$ be the dual rotation of the rotation containing (a, b) (if there is a stable table exposing a nonsingular rotation containing (a, b)). Considering Example 1, we have e.g. $\rho^{m_1, w_2} = \varphi_2$.

Lemma 3 ([20, Corollary 5.1]). *Let $\{a, b\}$ be a stable pair such that there is a stable pair $\{a, b'\}$ with a preferring b to b' . Then, there is a rotation including (a, b) . Moreover, $\rho^{a,b}$ is the unique rotation whose elimination makes b the last choice of a .*

Lemma 3 directly implies that in case a closed and complete subset Z contains $\rho^{a,b}$, agent a cannot be matched worse than b in the matching corresponding to Z :

Lemma 4. *Let $\{a, b\}$ be a stable pair such that there is a stable pair $\{a, b'\}$ with a preferring b to b' and let M be the stable matching corresponding to a closed and complete subset Z . If $\rho^{a,b} \in Z$, then $\{a, b\} \in M$ or a prefers $M(a)$ to b .*

PROOF. If we eliminate $\rho^{a,b}$, then by Lemma 3, agent b will become last in the preferences of a . Thus, a needs to be matched to b or better in the resulting matching. \square

Combining Lemmas 2 and 3 gives a characterization of when a pair $\{a, b\}$ is contained in a stable matching:⁸

Lemma 5. *Let $\{a, b\}$ be a stable pair such that there is a stable pair $\{a, b'\}$ with a preferring b to b' and let M be the stable matching corresponding to a closed and complete subset Z . Then $\{a, b\} \in M$ if and only if $\rho^{a,b} \in Z$ and for any stable partner b^* which a prefers to b , we have $\rho^{a,b^*} \notin Z$.*

PROOF. We start by proving the forward direction. Let M be a stable matching with $\{a, b\} \in M$ corresponding to the closed and complete subset Z of rotations. Successively eliminating rotations from Z to arrive at matching M , at some point b needs to become the last choice of a . By Lemma 3 for this we need to eliminate rotation $\rho^{a,b}$, implying that $\rho^{a,b} \in Z$. Moreover, note that in case we eliminate a rotation ρ^{a,b^*} where b^* is a stable partner of a which a prefers to b , then by Lemma 3 agent b^* becomes the last agent in the preferences of a . As a prefers b^* to b , this implies that b got deleted from the preferences of a , a contradiction.

For the backwards direction, assume that $\rho^{a,b} \in Z$ and $\rho^{a,b^*} \notin Z$ for every stable partner b^* of a which a prefers to b . As $\rho^{a,b} \in Z$, Lemma 4 implies that a is matched at least as good as b in M . Assume for the sake of contradiction that a is matched to an agent b^* it prefers to b in M . However, for b^* to become the only agent in

⁸Lemma 5 has been already implicitly used in the literature, e.g., [16, Section 5], but we are not aware of an explicit formulation or proof of it.

the preferences of a it in particular needs to become the last agent. By Lemma 3 this requires $\rho^{a,b^*} \in Z$, a contradiction. \square

Going back to our initially described intuition, for stable pairs $\{a, b\}$ covered by Lemma 5, $\rho^{a,b}$ can be interpreted as the necessary rotation and the rotations ρ^{a,b^*} for all stable partners b^* which a prefers to b can be interpreted as the prohibited rotations. To give an example for this, consider again Example 1, and let us focus on the stable pair $\{m_1, w_2\}$. Agent m_1 has a stable partner w_1 it prefers to w_2 and a stable partner w_3 to which it prefers w_2 . Thus, by Lemma 5, for $\{m_1, w_2\}$ to be included in a stable matching, the corresponding set of rotations needs to include $\rho^{m_1, w_2} = \varphi_2$ and cannot include $\rho^{m_1, w_1} = \varphi_4$ (in fact the single stable matching containing $\{m_1, w_2\}$ corresponds to the rotation set $\{\varphi_1, \varphi_2\}$).

Finally, we conclude by observing that for every stable pair $\{a, b\}$ and each stable matching M not including $\{a, b\}$ exactly one of a and b prefers the other to its partner in M :

Lemma 6 ([21, Lemma 4.3.9]). *Let M be a stable matching and $e = \{a, b\} \notin M$ be a stable pair. Then either $M(a) \succ_a b$ and $a \succ_b M(b)$ or $b \succ_a M(a)$ and $M(b) \succ_b a$.*

4.2.3 The Algorithm. Using the machinery from Section 4.2.2, we are now ready to present our algorithm.

Theorem 2. *ADAPT SR TO FORCED AND FORBIDDEN PAIRS can be solved in $O(2^{|P \cap M_1|} \cdot n \cdot m)$ time.*

PROOF. In the algorithm, we will guess⁹ for each forbidden pair $e = \{a, b\} \in P \cap M_1$ whether a or b prefers its partner in the output matching to its partner in M_1 . We say that a matching M respects our guesses if for each forbidden pair $e = \{a, b\} \in P \cap M_1$, a prefers its partner in M to b if and only if we guessed that this is the case. We assume that there is at least one stable matching containing all forced and none of the forbidden pairs that respects our guesses, as we can reject the current guess otherwise (and this can be checked in $O(m)$ time by reducing it to an instance of STABLE ROOMMATES WITH FORCED AND FORBIDDEN PAIRS [16]). We further assume without loss of generality that P only contains stable pairs (otherwise, we can delete the pair from P , as each stable matching will trivially not contain this pair).

In the following, when we say that we integrate a (nonsingular) rotation φ in a closed and complete set Z of rotations, then we add φ and all rotations preceding φ to Z and delete $\bar{\varphi}$ and all rotations preceded by $\bar{\varphi}$ from Z . Before we present the algorithm, we now argue that after integrating a nonsingular rotation φ to a closed and complete set Z , the resulting set Z' is still closed and complete: Z' is closed, as Z is closed and in case we add a rotation we also add all its predecessors and in case we delete a rotation we also delete all its successors. Moreover, Z' is complete: When integrating φ , we first add φ and delete $\bar{\varphi}$. For all other rotations that we add, i.e., all rotations preceding φ , we delete their dual and for all “dual” rotations we delete, i.e., all rotations succeeding $\bar{\varphi}$, we add the “primal”, as $\varphi \triangleright \rho$ if and only if $\bar{\rho} \triangleright \bar{\varphi}$ [21, Lemma 4.3.7].

In the algorithm, we start with a closed and complete subset of rotations Z and then only modify Z by integrating rotations. Thus, Z remains to be closed and complete over the course of the algorithm. We denote as M_Z the stable matching corresponding

⁹“Guessing” can be interpreted as iterating over all possibilities.

to Z (the correspondence between matchings and sets of rotations is described in Lemma 2).

The Algorithm. Our algorithm works as follows:

- (1) Compute the rotation digraph which contains a vertex for each rotation and an arc from rotation φ to rotation ρ if φ precedes ρ . Let Z_1 be the closed complete subset of rotations corresponding to M_1 , which exists and is unique by Lemma 2. Set $Z := Z_1$.
- (2) For each forced pair $\{a, b\} \in Q$ that is not a fixed pair, assume without loss of generality that there is a stable pair $\{a, b'\}$ with a preferring b to b' (for one of the two agents such a pair needs to exist by Lemma 6 and as $\{a, b\}$ is stable but not fixed). We integrate rotation $\rho^{a,b}$ to Z . Further, for each stable pair $\{a, b^*\}$ with a preferring b^* to b , we integrate $\bar{\rho}^{a,b^*}$ to Z .
- (3) For each forbidden pair $e = \{a, b\} \in P \cap M_1$, we guess whether a or b prefers its partner in the desired matching to its partner in M_1 (note that by Lemma 6, exactly one of a and b has to do this). We assume without loss of generality that we guessed that a prefers its partner in the desired matching to b . Let b^* be the least-preferred (by a) stable partner of a which a prefers to b (such a partner needs to exist by our guess). Integrate ρ^{a,b^*} to Z .
- (4) As long as the matching M_Z contains a pair $e = \{a, b\} \in P \setminus M_1$, assume without loss of generality that a prefers $M_Z(a) = b$ to $M_1(a)$ (for one of the two agents this needs to hold by Lemma 6, as $\{a, b\}$ is a stable pair). Let b^* be the least-preferred (by a) stable partner of a which a prefers to b . If b^* exists, we integrate ρ^{a,b^*} to Z ; otherwise we do nothing.
- (5) Return the matching $M := M_Z$.

Proof of Correctness. We start by showing that all changes made to Z over the course of the algorithm are indeed necessary.

Claim 1. *Let M^* be a stable matching containing all forced pairs and no forbidden pairs which respects our guesses. Further, let Z^* be the corresponding closed and complete subset of rotations. Then Z^* contains all rotations added in Steps 2 to 4. Moreover, the agent b^* defined in Step 4 always exists.*

PROOF OF CLAIM. We start by proving the first part of the claim. Note that it is sufficient to prove the statement for all integrated rotations, as all other rotations ρ that we added to Z precede an integrated rotation.

Each rotation integrated in Step 2 is contained in Z^* by Lemma 5 and as Z^* needs to be complete.

Next, we consider the rotations integrated in Step 3. Assume without loss of generality that we guessed that a prefers its partner in the desired matching to b . Then a must be matched at least as good as b^* in the desired matching. Assume towards a contradiction that $\rho^{a,b^*} \notin Z^*$, implying $\bar{\rho}^{a,b^*} \in Z^*$. Rotation $\bar{\rho}^{a,b^*}$ contains (a, b^*) (by definition of $\bar{\rho}^{a,b^*}$ and as the dual of $\bar{\rho}$ is again ρ). Recall that in case $\bar{\rho}^{a,b^*}$ is exposed, then a is the last choice of b^* (Lemma 1) and eliminating the rotation implies deleting the pair $\{a, b^*\}$ (Lemma 1). Thus, after the elimination of $\bar{\rho}^{a,b^*}$, agent b^* prefers its last choice to a . This implies that b^* prefers $M^*(b^*)$ to a . By Lemma 6 and as $\{a, b^*\}$ is a stable pair not contained in M^* , it follows that a

prefers b^* to $M^*(a)$. This is a contradiction to a preferring $M^*(a)$ to b and the definition of b^* . Consequently, $\rho^{a,b^*} \in Z^*$.

For Step 4, we show by induction that the claim holds after the i -th execution of this step. The statement clearly holds before the first execution of the step. Let Z^i be the set Z before the i -th execution. Let $\{a, b\}$ be the pair examined in this execution, with a preferring b to $M_1(a)$. Lemma 5 implies that $\rho^{a,b} \in Z^i$, as a is matched to b in M_{Z^i} and $M_1(a)$ is a stable partner of a to which a prefers b . Moreover, we need to have that $\rho^{a,b} \notin Z_1$: If $\rho^{a,b} \in Z_1$, then by Lemma 4, a is matched at least as good as b in M_1 , contradicting our assumption that a prefers b to $M_1(a)$. By our induction hypothesis it follows that $\rho^{a,b} \in Z^*$. Applying again Lemma 4 it follows that a is matched at least as good as b in M^* . As $\{a, b\}$ is a forbidden pair, we even get that a prefers $M^*(a)$ to b . The remainder of the proof is now analogous to Step 3.

Concerning the second part of the claim, observe that we have established above that in each iteration of Step 4, a prefers $M^*(a)$ to b . From this it follows that a has a stable partner it prefers to b and thus in particular that b^* exists in each execution of Step 4. ■

Recall that we have assumed that there is a stable matching M^* containing all forced and no forbidden pairs that respects our guesses. Thus, as all rotations integrated to Z must be contained in Z^* by Claim 1, there is no rotation ρ such that ρ as well as $\bar{\rho}$ get added to Z during the algorithm (as in this case, Z^* would not be a complete subset of rotations). Step 2 now ensures by Lemma 5 that M contains all forced pairs. Steps 3 and 4 ensure that M contains no forbidden pair by Lemma 4 (note that the case that b^* does not exist in Step 4 never occurs as proven in Claim 1).

Next, we show the optimality of M . Let Z^* be the subset of the rotation poset corresponding to M^* . By Claim 1, we get $Z_1 \Delta Z \subseteq Z_1 \Delta Z^*$ (Claim 1 directly implies that $Z \setminus Z_1 \subseteq Z^* \setminus Z_1$ but also gives us $Z_1 \setminus Z \subseteq Z_1 \setminus Z^*$ as deleting a rotation corresponds to adding its dual). We now show that we can conclude from this that there is no pair $e \in (M_1 \cap M^*) \setminus (M_1 \cap M)$: Assume towards a contradiction that there is some $e = \{a, b\} \in (M_1 \cap M^*) \setminus (M_1 \cap M)$. Note that as $\{a, b\}$ is not contained in the stable matching M , it is not a fixed pair. Assume without loss of generality that there is a stable pair $\{a, b'\}$ with a preferring b to b' (for one of the two agents this needs to exist by Lemma 6, as $\{a, b\}$ is a stable pair not contained in M). Thus, by Lemma 5, $Z_1 \cap Z^*$ contain $\rho^{a,b}$ as well as $\bar{\rho}^{a,b^*}$ for any stable partner b^* which a prefers to b . Since $Z_1 \cap Z^* \subseteq Z_1 \cap Z$ these rotations are also contained in Z and by Lemma 5 it follows that $\{a, b\}$ is also contained in M , a contradiction to $\{a, b\} \in (M_1 \cap M^*) \setminus (M_1 \cap M)$.

Running Time. Computing the rotation digraph can be done in $O(n \cdot m)$ time [14]. In Step 3, there are $2^{|P \cap M_1|}$ guesses. For each guess, any pair can be added at most once to M_Z and any rotation can be added at most once to Z . Thus, the remaining part of Steps 2 to 4 can be done in $O(m)$ total time. Consequently, the algorithm runs in $O((2^{|P \cap M_1|} + n) \cdot m)$ time. □

4.3 (FPT-)Algorithm for Adapt Strongly SR with Ties to Forced and Forbidden Pairs

In case of strong stability, we can employ a similar algorithm as for strict preferences, as this problem also admits a (slightly more

complicated) rotation poset. Although the definition of the rotations and their duals differ from the “classical” case without ties, they still fulfill crucial properties exploited in Theorem 2:

- (1) Analogous to Lemma 2, each stable matching corresponds to a closed and complete subset of the rotation poset.
- (2) Somewhat analogous to Lemma 5, for each stable pair e there are two rotations ρ and φ such that e may be contained in a stable matching corresponding to a complete and closed set Z of rotations if and only if $\rho \in Z$ and $\varphi \notin Z$.

There are now multiple possible stable matchings corresponding to the same set of rotations (as rotations here only encode the rank of the partner of an agent in the matching). In order to solve ADAPT STRONGLY SR WITH TIES TO FORCED AND FORBIDDEN PAIRS, it suffices to compute the closed and complete set Z of rotations corresponding to an optimal solution (as subsequently we can find the stable matching corresponding to Z closest to M_1 using a minimum-cost matching algorithm). Turning to the constraints that forced and forbidden pairs impose on Z , as for strict preferences, each forced pair gives rise to the constraint that one rotation is contained and one rotation is not contained in Z (due to part (2) of the above enumeration). For forbidden pairs the situation is different and more complicated: As there may be multiple stable matchings for the same set of rotations with only some of them not containing a forbidden pair in question, a forbidden pair does not necessarily lead to a constraint on Z (even if this forbidden pair is also contained in M_1). In order to be able to solve the problem, we show as a crucial step that we can determine whether (a set of) forbidden pairs lead to constraints on Z . We can show the following:

Theorem 3 (★). *ADAPT STRONGLY SR WITH TIES TO FORCED AND FORBIDDEN EDGES can be solved in $O((2^{|P \cap M_1|} + m) \cdot \sqrt{nm} \log n)$ time.*

5 CONCLUSION

We have conducted a complete and fine-grained analysis of minimally changing stable matchings to include forced and exclude forbidden pairs. As our main result, we have proven that ADAPT SR TO FORCED AND FORBIDDEN PAIRS is fixed-parameter tractable with respect to the number of forbidden pairs in the given matching (and thus polynomial-time solvable if there are only forced pairs). At the core of this algorithm lies a clever exploitation of the rotation poset that might inspire similar approaches for related problems. For example, one might want to adapt a matching to (dis)satisfy certain groups of agents or to improve the situation of the agent which is worst of. All these requirements can be encoded if we are given for each agent an upper and lower bound for how the agent is matched in the new matching M_2 (i.e., for each agent $a \in A$, we are given two agents b_a and v_a and we require $b_a >_a M_2(a) >_a v_a$). Slightly adapting the initialization procedure of our algorithm (by starting with incorporating rotations realizing these constraints) this problem becomes polynomial-time solvable.

To the best of our knowledge, the idea of minimally changing a given matching to incorporate external requirements has not been studied in previous works. Thus, extending our studies of forced and forbidden pairs for stable matchings to other incremental requirements such as group fairness or diversity constraints or other matching problems such as popular matching is an interesting direction for future work.

Acknowledgments

NB was supported by the DFG project ComSoc-MPMS (NI 369/22). KH was supported by the DFG project FPTinP (NI 369/16).

REFERENCES

- [1] Sayan Bhattacharya, Martin Hoefer, Chien-Chung Huang, Telikepalli Kavitha, and Lisa Wagner. 2015. Maintaining Near-Popular Matchings. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP '15)*. Springer, 504–515.
- [2] Niclas Boehmer and Klaus Heeger. 2022. Adapting Stable Matchings to Forced and Forbidden Pairs. *CoRR* abs/2204.10040 (2022). arXiv:2204.10040 <https://arxiv.org/abs/2204.10040>
- [3] Niclas Boehmer, Klaus Heeger, and Rolf Niedermeier. 2022. Deepening the (Parameterized) Complexity Analysis of Incremental Stable Matching Problems. In *Proceedings of the 47th International Symposium on Mathematical Foundations of Computer Science (MFCS '22)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 21:1–21:15.
- [4] Niclas Boehmer, Klaus Heeger, and Rolf Niedermeier. 2022. Theory of and Experiments on Minimally Invasive Stability Preservation in Changing Two-Sided Matching Markets. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI '22)*. AAAI Press, 4851–4858.
- [5] Niclas Boehmer and Rolf Niedermeier. 2021. Broadening the Research Agenda for Computational Social Choice: Multiple Preference Profiles and Multiple Solutions. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '21)*. ACM, 1–5.
- [6] Robert Brederbeck, Jiehua Chen, Dušan Knop, Junjie Luo, and Rolf Niedermeier. 2020. Adapting Stable Matchings to Evolving Preferences. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI '20)*. AAAI Press, 1830–1837.
- [7] Katarína Cechlářová, Laurent Gourvès, and Julien Lesca. 2019. On the Problem of Assigning PhD Grants. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI '19)*. ijcai.org, 130–136.
- [8] Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. 2004. Incremental Clustering and Dynamic Information Retrieval. *SIAM J. Comput.* 33, 6 (2004), 1417–1440.
- [9] Ágnes Cseh and Klaus Heeger. 2020. The stable marriage problem with ties and restricted edges. *Discret. Optim.* 36 (2020), 100571.
- [10] Ágnes Cseh and David F. Manlove. 2016. Stable Marriage and Roommates problems with restricted edges: Complexity and approximability. *Discret. Optim.* 20 (2016), 62–89.
- [11] Vânia M. F. Dias, Guilherme Dias da Fonseca, Celina M. H. de Figueiredo, and Jayme Luiz Szwarcfiter. 2003. The stable marriage problem with restricted pairs. *Theor. Comput. Sci.* 306, 1-3 (2003), 391–405.
- [12] David Eisenstat, Claire Mathieu, and Nicolas Schabanel. 2014. Facility Location in Evolving Metrics. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP '14)*. Springer, 459–470.
- [13] Tomás Feder. 1992. A New Fixed Point Approach for Stable Networks and Stable Marriages. *J. Comput. Syst. Sci.* 45, 2 (1992), 233–284.
- [14] Tomás Feder. 1994. Network Flow and 2-Satisfiability. *Algorithmica* 11, 3 (1994), 291–319.
- [15] Itai Feigenbaum, Yash Kanoria, Irene Lo, and Jay Sethuraman. 2020. Dynamic Matching in School Choice: Efficient Seat Reallocation After Late Cancellations. *Manag. Sci.* 66, 11 (2020), 5341–5361.
- [16] Tamás Fleiner, Robert W. Irving, and David F. Manlove. 2007. Efficient algorithms for generalized Stable Marriage and Roommates problems. *Theor. Comput. Sci.* 381, 1-3 (2007), 162–176.
- [17] Karthik Gajulapalli, James A. Liu, Tung Mai, and Vijay V. Vazirani. 2020. Stability-Preserving, Time-Efficient Mechanisms for School Choice in Two Rounds. In *Proceedings of the 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS '20)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 21:1–21:15.
- [18] D. Gale and L. S. Shapley. 2013. College Admissions and the Stability of Marriage. *Am. Math. Mon.* 120, 5 (2013), 386–391.
- [19] Sushmita Gupta, Pallavi Jain, Sanjukta Roy, Saket Saurabh, and Meirav Zehavi. 2020. On the (Parameterized) Complexity of Almost Stable Marriage. In *Proceedings of the 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS '20)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 24:1–24:17.
- [20] Dan Gusfield. 1988. The Structure of the Stable Roommate Problem: Efficient Representation and Enumeration of All Stable Assignments. *SIAM J. Comput.* 17, 4 (1988), 742–769.
- [21] Dan Gusfield and Robert W. Irving. 1989. *The Stable Marriage Problem – Structure and Algorithms*. MIT Press.
- [22] Guillaume Haeringer and Vincent Iehlé. 2021. Gradual college admission. *J. Econ. Theory* 198 (2021), 105378.
- [23] Kathrin Hanauer, Monika Henzinger, and Christian Schulz. 2022. Recent Advances in Fully Dynamic Graph Algorithms (Invited Talk). In *Proceedings of the 1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND '22)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 1:1–1:47.
- [24] Robert W. Irving. 1985. An Efficient Algorithm for the "Stable Roommates" Problem. *J. Algorithms* 6, 4 (1985), 577–595.
- [25] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA (The IBM Research Symposia Series)*. Plenum Press, New York, 85–103.
- [26] Donald E. Knuth. 1976. *Mariages stables et leurs relations avec d'autres problèmes combinatoires*. Les Presses de l'Université de Montréal, Montréal, Que. 106 pages. Introduction à l'analyse mathématique des algorithmes, Collection de la Chaire Aisenstadt.
- [27] Adam Kunysz. 2018. An Algorithm for the Maximum Weight Strongly Stable Matching Problem. In *Proceedings of the 29th International Symposium on Algorithms and Computation (ISAAC '18)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 42:1–42:13.
- [28] David F. Manlove. 1999. *Stable marriage with ties and unacceptable partners*. Technical Report. University of Glasgow, Department of Computing Science.
- [29] David F. Manlove. 2013. *Algorithmics of Matching Under Preferences*. Series on Theoretical Computer Science, Vol. 2. WorldScientific.
- [30] David F. Manlove, Robert W. Irving, Kazuo Iwama, Shuichi Miyazaki, and Yasufumi Morita. 2002. Hard variants of stable marriage. *Theor. Comput. Sci.* 276, 1-2 (2002), 261–279.
- [31] Dániel Marx and Ildikó Schlotter. 2010. Parameterized Complexity and Local Search Approaches for the Stable Marriage Problem with Ties. *Algorithmica* 58, 1 (2010), 170–187.
- [32] Alvin Roth. 1986. On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica* (1986), 425–427.