

Adapting Stable Matchings to Evolving Preferences

Robert Brederick,¹ Jiehua Chen,² Dušan Knop,^{1,3} Junjie Luo,¹ Rolf Niedermeier¹

¹Technische Universität Berlin, Chair of Algorithmics and Computational Complexity

²Algorithms and Complexity Group, TU Wien, Vienna, Austria

³Department of Theoretical Computer Science, Faculty of Information Technology,
Czech Technical University in Prague, Prague, Czech Republic

{robert.bredereck, junjie.luo, rolf.niedermeier}@tu-berlin.de, jiehua.chen@tuwien.ac.at, dusan.knop@fit.cvut.cz

Abstract

Adaptivity to changing environments and constraints is key to success in modern society. We address this by proposing “incrementalized versions” of STABLE MARRIAGE and STABLE ROOMMATES. That is, we try to answer the following question: for both problems, what is the computational cost of adapting an existing stable matching after some of the preferences of the agents have changed. While doing so, we also model the constraint that the new stable matching shall be not too different from the old one. After formalizing these incremental versions, we provide a fairly comprehensive picture of the computational complexity landscape of INCREMENTAL STABLE MARRIAGE and INCREMENTAL STABLE ROOMMATES. To this end, we exploit the parameters “degree of change” both in the input (difference between old and new preference profile) and in the output (difference between old and new stable matching). We obtain both hardness and tractability results, in particular showing a fixed-parameter tractability result with respect to the parameter “distance between old and new stable matching”.

Introduction

Imagine the following scenario. A manager responsible for a group of $2n$ workers has to form n two-worker teams based on the preferences over potential work partners of each worker. The manager, being interested in a robust solution, computes a stable matching (indeed, this refers to the STABLE ROOMMATES problem). However, say every month the workers may update their preferences about wanted work partners (the updates may be based on gained experiences, new information, newly developed personal skills etc.). The team manager then has to find a new stable matching respecting the individually evolved preferences. To this end, however, the team manager may not want to allow too radical changes in the composition of the two-worker teams because this might e.g. overburden administration. Thus, a moderate change is acceptable, but too radical changes in the team compositions should be avoided whenever possible. We address this scenario by introducing and studying “incremental versions” of the two most prominent stable matching problems, namely STABLE MARRIAGE and STABLE ROOMMATES.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In stable matching scenarios or, in other words, in matching under preferences (Manlove 2013), one is given a set of agents, each of them having preferences over (some of) the other agents, and the goal is to match pairs of agents such that the outcome is stable. Informally speaking, stability means that there are no two agents that would both prefer to be matched with each other instead to their current partners in the matching (or being unmatched). Two classic problems here are the bipartite case with two equal-sized sets of agents (referred to as STABLE MARRIAGE) and the general case of an even number of agents (referred to as STABLE ROOMMATES). Motivated by the introductory considerations, we next define “incremental versions” of both problems; further formal definitions are presented later.

INCREMENTAL STABLE MARRIAGE (resp. INCREMENTAL STABLE ROOMMATES)

Input: Two disjoint sets U and W of n agents each (resp. a set V of $2n$ agents), two preference profiles \mathcal{P}_1 and \mathcal{P}_2 for $U \uplus W$ (resp. V), a stable matching M_1 for profile \mathcal{P}_1 , and a non-negative integer k .

Question: Does $U \uplus W$ (resp. V) admit a stable matching M_2 for profile \mathcal{P}_2 such that $\text{dist}(M_1, M_2) = |M_1 \Delta M_2| \leq k$?

Herein, $M_1 \Delta M_2$ is the symmetric difference between the two matchings (sets of edges) M_1 and M_2 .

In both definitions there are two main regulating screws. First, we are given *two* preference profiles, the old \mathcal{P}_1 and the new \mathcal{P}_2 , thereby reflecting the change of preferences. Indeed, to reflect only moderate changes (“evolution”), we will subsequently measure the difference between the two profiles (later referred to as swap distance), yielding a natural problem-specific parameter (the smaller it is, the less revolutionary the changes are). Second, the number k can be interpreted as a locality parameter—it exposes how close the new matching has to be to the old one. The smaller we choose k , the more conservative we are with respect to change in the outcome (namely the difference between old and new stable matching). Together, we thus have one parameter to regulate the degree of change measured in the *input* preferences and one parameter to regulate the degree of change measured in the *output* stable matching solution. Taking up these two parameters, we provide a thorough parameterized complexity analysis of

these kinds of stable matching problems with evolving preferences. To this end, we also distinguish between preferences with and without ties. Roughly speaking, we provide positive (fixed-parameter) tractability results in the case without ties and several (parameterized) intractability results for the case with ties. Before describing our results in more detail, however, we discuss related work.

Related work. There is previous work on matching-related problems in the context of dynamic graph algorithms where vertices and/or edges arrive or depart iteratively over time. The goal then is to maintain a solution of sufficient quality by performing necessary updates after every single change. The main difference to our work is that we study changes between two preference profiles (not so much in the graph structure), and the changes can be at a larger scale. For instance, Bhattacharya et al. (2015) studied the maintenance of near-popular matchings (a scenario related to stable matchings) based on a greedy improvement strategy, heavily employing maximum vertex degree in their analysis. Ghosal, Kunysz, and Paluch (2017) and Nimbhorkar and Rameshwar (2019) investigated dynamic rank-maximal and popular matchings, again within the setting of dynamic graph algorithms and a focus on update times after each single change in the graph. Kanade, Leonardos, and Magniez (2016) studied STABLE MARRIAGE where at each time-step, two random adjacent agents in some preference list are swapped, and designed approximation algorithms to maintain a matching with logarithmic number of blocking pairs.

Genc et al. (2017a; 2017b; 2019) studied *robustness* of a matching in STABLE MARRIAGE. Herein, the robustness of a given stable matching is measured by the number of modifications needed to find an alternative stable matching if some currently matched agent pairs break up. They define an (x, y) -supermatch as a stable matching that satisfies the following: If any x agents break up, then it is possible to rematch these x agents so that the new matching is again stable; further, this re-matching *must not* break up more than y other pairs. Genc et al. (2019) showed that deciding whether a $(1, 1)$ -supermatch exists is NP-complete. The main distinguishing features compared to our model are that, using our two regulating screws mentioned above, we can model both moderate changes in the preferences (that is, the input) and moderate differences between the old and the new stable matching (that is, the output). Indeed, we perform a parameterized complexity analysis exploiting these parameters while Genc et al. (2017a; 2017b; 2019) focused on classic complexity results and used heuristics in experimental work. Moreover, we model changing preferences while they addressed breaking up matched pairs. Finally, our main contributions are in the STABLE ROOMMATES case while they exclusively focused on STABLE MARRIAGE without ties.

Our model of measuring the distance between the input and sought matching is related to the stable matching problem with *forbidden* and *forced* edges studied by Cseh and Manlove (2016). The goal is to find a stable matching which minimizes the number of forbidden edges plus the number of non-forced edges. While we do not model forbidden edges,

Cseh and Manlove do not assume changes in agents' preferences. This makes a difference in terms of parameterized complexity.

Marx and Schlotter (2010; 2011) studied local search aspects for the NP-hard STABLE MARRIAGE with ties. More specifically, they investigated the parameterized complexity of a local search variant of STABLE MARRIAGE using parameters such as the number of ties. Thus, the main overlap with our work is in terms of searching for local improvements and employing parameterized complexity analysis—the studied computational problems are different from ours as they do not model changes in the input preferences.

There have been numerous other models and investigations to enrich the basic stable matching model, including the use of only partially ordered preferences (Drummond and Boutilier 2013), “multilayer” stable matchings with several preference profiles to be obeyed ‘in parallel’ (Aziz et al. 2016; Miyazaki and Okamoto 2017; Chen, Niedermeier, and Skowron 2018), or studying robust stability in a probabilistic model (Mai and Vazirani 2018a; 2018b) or from a quantitative angle (Chen, Skowron, and Sorge 2019).

Finally, let us briefly mention that the motivation for our incremental scenario for stable matching is related to similar scenarios in the context of clustering (Charikar et al. 2004; Luo et al. 2018), coloring (Hartung and Niedermeier 2013), other dynamic versions of parameterized problems (Abu-Khzam et al. 2015; Krithika, Sahu, and Tale 2018), and reoptimization (Böckenhauer et al. 2018; Schieber et al. 2018).

Our contributions. Besides introducing a fresh model of stable matching computations, we provide results mostly in terms of parameterized complexity analysis for both INCREMENTAL STABLE MARRIAGE and INCREMENTAL STABLE ROOMMATES, where we see the main technical contributions mostly for the latter. In particular, our main algorithmic result is that INCREMENTAL STABLE ROOMMATES for input instances without ties is fixed-parameter tractable with respect to the parameter k (distance between the old and the new matchings). To show this, we heavily use structural results due to Irving (1994) and Gusfield (1988) and show how to exploit them for designing a fixed-parameter algorithm.

Most of our results are surveyed in Table 1. Herein, $\mathcal{P}_1 \oplus \mathcal{P}_2$ denotes the swap distance between two preference profiles (see the next section for formal definitions). The table indicates that we obtained a fairly complete picture of the computational complexity landscape, e.g., also complementing $W[1]$ -hardness results with corresponding XP-algorithms.

Finally, we mention in passing that to achieve our results, throughout the work we also introduce and study some intermediate problems (for instance, EDGE-INCREMENTAL INDEPENDENT SET) which may be of independent interest and prove useful in other settings.

Due to space constraints, many technical details and proofs are deferred to the full version (Bredereck et al. 2019).

Definitions and notations

In this section, we review fundamental concepts used in matchings under preferences.

ties	$ \mathcal{P}_1 \oplus \mathcal{P}_2 $	INCR. ST. MARRIAGE		INCR. STABLE ROOMMATES	
no	any	P	(Prop. 1)	FPT for $k := \text{dist}(M_1, M_2)$, W[1]-h for $k' := M_1 \cap M_2 $	(Thm. 1), (Prop. 5)
yes	1	W[1]-h for k	(Thm. 2)	W[1]-h for k , even for compl. pref., NP-h even if $k' \geq 0$	(Thm. 4)
yes	2	W[1]-h for k'	(Thm. 3)	W[1]-h for k'	(Thm. 3)
yes	any	XP for k	(Prop. 2)	XP for k	(Prop. 2)

Table 1: Overview of our results. Unless otherwise stated, results are for the general case where preferences could be incomplete.

Preference lists and profiles. Let $V = \{1, 2, \dots, 2n\}$ be a set of $2n$ agents. Each agent $i \in V$ has a subset $V_i \subseteq V \setminus \{i\}$ of agents which they find *acceptable* as partners and has a *preference list* \succeq_i on V_i (i.e., a transitive and complete binary relation on V_i). Here, $x \succeq_i y$ means that i weakly prefers x over y (i.e., x is at least as good as y). We use \succ_i to denote the asymmetric part (i.e., $x \succeq_i y$ and $\neg(y \succeq_i x)$) and \sim_i to denote the symmetric part of \succeq_i (i.e., $x \succeq_i y$ and $y \succeq_i x$). If $x \sim_i y$, then we also say that x and y are *tied* in i 's preference list and that agent i has *ties* in its preference list.

For two disjoint subsets of agents $X \subseteq V$ and $Y \subseteq V$ we write $X \succeq_i Y$ if for each pair of agents $x \in X$ and $y \in Y$ we have $x \succeq_i y$.

A preference profile \mathcal{P} for V is a collection $(\succeq_i)_{i \in V}$ of preference lists for each agent $i \in V$. A profile \mathcal{P} may have the following properties:

1. It is *complete* if for each agent $i \in V$ it holds that $V_i \cup \{i\} = V$; otherwise it is *incomplete*.
2. The profile \mathcal{P} has *ties* if there is an agent $i \in V$ with ties in its preference list.

To an instance (V, \mathcal{P}) we assign an *acceptability graph* G , which has V as its vertex set and two agents are connected by an edge if each finds the other acceptable. Without loss of generality, G does not contain isolated vertices, meaning that each agent has at least one agent which it finds acceptable.

Swaps and differences between two matchings. Given two preference lists \succeq and \succeq' , the *swap distance between \succeq and \succeq'* is defined as the number of pairs that are “ordered” differently; if \succeq and \succeq' are defined on different sets, then we assume that their swap distance is infinite. Formally, $\delta(\succeq, \succeq') := \infty$ if \succeq and \succeq' are defined on *different* sets, otherwise, $\delta(\succeq, \succeq') := |\{(x, y) \mid x \succ y \wedge y \succeq' x\}| + |\{(x, y) \mid x \sim y \wedge (x, y) \notin \sim'\}|$. For example, $\delta(x \sim y \sim z, z \succ \{x, y\}) = 2$.

Let \mathcal{P}_1 and \mathcal{P}_2 be two preference profiles for the same set V of agents. The *swap distance between \mathcal{P}_1 and \mathcal{P}_2* , denoted as $|\mathcal{P}_1 \oplus \mathcal{P}_2|$, is defined as the sum of the swap distances between the preference lists of each agent in the two preference profiles. Formally, let \succeq_i^j be the preference list of an agent $i \in V$ in the profile \mathcal{P}_j for $j = 1, 2$. We have that $|\mathcal{P}_1 \oplus \mathcal{P}_2| = \sum_{i \in V} \delta(\succeq_i^1, \succeq_i^2)$.

Given a set V of agents, a *matching* M of V is a set of pairwise disjoint pairs of agents in V . Given two matchings M_1 and M_2 for the same set V with $2n$ agents, we define the *distance between M_1 and M_2* as the size of the symmetric difference of M_1 and M_2 , formally: $\text{dist}(M_1, M_2) := |M_1 \Delta M_2| = |M_1 \setminus M_2| + |M_2 \setminus M_1|$.

Blocking pairs and stable matchings. Let a preference profile \mathcal{P} for a set V of agents, the corresponding acceptability graph G , and a matching $M \subseteq E(G)$ be given. For a pair $\{x, y\}$ of agents, if $\{x, y\} \in M$, then we denote the corresponding partner y by $M(x)$; otherwise we call this pair *unmatched*. We write $M(x) = \perp$ if agent x has *no partner*, i.e., if the agent x is not involved in any pair in M . We use $\perp(M)$ to denote the set of unmatched agents in a matching M , that is, $\perp(M) = \{x \mid M(x) = \perp\}$. If *no* agent x has $M(x) = \perp$ (i.e., $\perp(M) = \emptyset$), then M is called *perfect*.

An unmatched pair $\{x, y\} \in E(G) \setminus M$ is *blocking* M if both x and y prefer each other to being unmatched or to their assigned partners, i.e., it holds that $(M(x) = \perp \vee y \succ_x M(x)) \wedge (M(y) = \perp \vee x \succ_y M(y))$. We call a matching M *stable*¹ if no unmatched pair is blocking M .

Algorithms

To start with, we note that our most restricted problem variant, INCREMENTAL STABLE MARRIAGE without ties, can be solved in polynomial time. The main idea behind this result is based on the fact that there exists a compact and polynomial-time computable representation of all stable matchings, the so-called partially ordered set of polynomially many rotations (Gusfield and Irving 1989, Chapter 2). Equipping these rotations with some appropriate weights, we can reduce our problem to finding a closed subset of rotations with maximum weight, which can be solved in polynomial time using an approach similar to a known one (Gusfield and Irving 1989, Chapter 3.6.1). We refer to the full version (Bredereck et al. 2019) for more details.

Proposition 1. INCREMENTAL STABLE MARRIAGE *without ties* can be solved in $O(n^3)$ time.

Second we show that our least restricted NP-hard problem variant, INCREMENTAL STABLE ROOMMATES with ties allowed, can at least be solved in polynomial time when the distance between the matchings M_1 and M_2 is a constant. In other words, Proposition 2 presents an XP-algorithm for the distance as a parameter.

Proposition 2. INCREMENTAL STABLE ROOMMATES *with ties* can be solved in $n^{O(k)}$ time, where $k = |M_1 \Delta M_2|$.

Proof. Let $\hat{n} = |M_1|$. The algorithm first guesses positive integers k_1, k_2 with $k_1 + k_2 \leq k$ which we further treat as

¹We exclusively focus on the (most common) weak stability concept (Manlove 2013). We conjecture that several results will also hold at least for strong stability. The proofs, however, may need non-trivial adjustments.

follows. We denote by k_1 the number of matching edges leaving M_1 (that is, $|M_1 \setminus M_2|$) and by k_2 the number of new edges (that is, $|M_2 \setminus M_1|$). Now, for each such pair (k_1, k_2) we first guess k_1 edges to delete from M_1 . Notice that there are $O(\hat{n}^{k_1}) = O(n^k)$ possible guesses. Let \hat{M}_1 be the rest of the matching M_1 (after we delete the just guessed edges). Then we guess k_2 pairs of vertices not matched in \hat{M}_1 . This completely defines the matching M_2 which can be checked in polynomial-time for stability. Notice that there are $\binom{2n-2\hat{n}+2k_1}{2k_2} \cdot \frac{(2k_2)!}{2 \cdot k_2!} = O(2k^{2k} \cdot (2n)^{2k})$ possible guesses. Repeating the above procedure for every pair (k_1, k_2) and applying the obvious bound $k \leq n$, we get the overall running time $n^{O(k)}$ as claimed. \square

We remark that Proposition 2 can also be shown by using an idea of Cseh and Manlove (2016, Theorem 4.14). Our approach, however, is much simpler. The most pressing question following from Proposition 2 is to ask for which cases this result can be improved to a fixed-parameter tractability result for the parameter k , the distance between M_1 and M_2 . Unfortunately, we will show in the next section that this is not possible when ties are involved. In the remainder of this section, however, we show that this is possible when ties are not allowed. Formally, we show the following.

Theorem 1. INCREMENTAL STABLE ROOMMATES *without ties can be solved in $O(2^k n^4)$ time.*

The algorithm behind Theorem 1 is partially inspired by the polynomial-time algorithm for solving MAXIMUM WEIGHT STABLE MARRIAGE (Gusfield and Irving 1989, Chapter 2.5.1). The high-level structure of our algorithm is as follows:

1. Using the algorithm of Irving (1985) and based on the structural insights of Gusfield (1988) that the set of all stable matchings can be compactly and efficiently represented via the so-called partially ordered set (poset) of polynomially many rotations, we compute the poset of rotations for the new preference profile P_2 .

Roughly speaking, a rotation (relative to a preference profile) involves a specific cyclic sequence ρ of agents where the first acceptable partner in the preference list of each agent x in ρ is exactly the second acceptable partner in the preference list of x 's predecessor in ρ . We remark that for a rotation removing the first acceptable partner of each agent in ρ will not alter the existence of a stable matching.

2. To capture the (different) costs of eliminating one rotation from each dual pair (see Definition 3) with respect to the resulting distance between the matchings M_1 and M_2 , we compute two weights for each rotation.
3. Finally, to consistently choose one rotation from each dual pair, and additionally respecting the weight constraints, we reduce our problem to an auxiliary graph problem defined below and provide a fixed-parameter algorithm for this auxiliary problem.

Each high-level step of the algorithm will be described in one of the subsequent subsections.

To define our auxiliary problem (which indeed may be of independent interest), for each partially ordered set (R, \succeq) , we say that a given subset $C \subseteq R$ is *closed with respect to the relation \succeq* if for each pair $\{\rho, \sigma\} \subseteq R$ of elements it holds that $\rho \in C$ and $\sigma \succeq \rho$ imply $\sigma \in C$.

WEIGHTED CONFLICT-FREE CLOSED SUBSET (WCS)

Input: A partially ordered set (R, \succeq) , an undirected graph $G = (R, E)$, a weight function $w: R \rightarrow \mathbb{N}$, a positive integer ℓ , and a budget $b \in \mathbb{N}$.

Question: Is there a closed subset $C \subseteq R$ of size ℓ which is independent in G such that $\sum_{c \in C} w(c) \leq b$?

Here, the conflicts are modeled by the edges so that the graph is also referred to as a *conflict graph* and seeking for an independent set solution means seeking for a *conflict-free set*.

Preprocessing and identification of the rotations

We first recall Irving's polynomial-time algorithm (Irving 1985) for determining whether an instance of STABLE ROOMMATES without ties has a stable matching and fundamental structural properties behind all stable matchings (Gusfield 1988). Irving's algorithm is divided into two phases: Phase 1 and Phase 2. Phase 1 involves a sequence of proposals from each agent i to the first agent j on i 's list, and each such proposal resulting in the deletion of all successors of i from j 's list. Phase 1 does not alter any stable matching since in this phase j is removed from i 's preference list (and i is removed from j 's preference list) only when $\{i, j\}$ does not form a pair in any stable matching. The set of lists at the end of the first phase is called *Phase 1 table*.

There are three possibilities for the Phase 1 table; note that the Phase 1 table is unique.

Every list in the table is empty. Then, there is no stable matching; Irving's original algorithm assumes that the input preferences are complete, implying that whenever a list becomes empty, then there is no stable matching; for incomplete preferences we can only infer the non-existence of stable matchings if every list becomes empty.

Every non-empty list contains exactly one agent. Then, we find a unique stable matching.

At least one list contains more than one agent. Then, we proceed to Phase 2.

For agents with empty preference lists in the Phase 1 table, we use the following.

Proposition 3 ((Gusfield and Irving 1989, Theorem 4.5.2)). *For each agent, if its preference list becomes empty after Phase 1, then it will not be matched by any stable matching; otherwise it must be matched by all stable matchings.*

Since in the first two cases the instance of INCREMENTAL STABLE ROOMMATES is trivial to solve, we assume the third case in the following and we can ignore every agent whose list becomes empty after Phase 1. Notice that in the Phase 1 table we always have that if an agent i is ranked first in j 's preference list, then j is ranked the last in i 's preference list. This invariant we keep throughout the whole Phase 2.

We need some definitions for Phase 2. A *preference table* is the Phase 1 table or any Phase 2 table (i.e., it is a collection of preference lists).

Definition 1 (Rotations). A *rotation* exposed in a preference table \mathcal{T} is an ordered sequence $(e_0, h_0), (e_1, h_1), \dots, (e_{r-1}, h_{r-1})$ of pairs of agents such that h_i and $h_{i+1 \bmod r}$ are the first and the second agent on e_i 's list in \mathcal{T} , for all $0 \leq i \leq r-1$.

Note that if (e, h) is a pair in a rotation, then e is ranked last by h with respect to \mathcal{T} .

Definition 2. The *elimination of an exposed rotation* $(e_0, h_0), (e_1, h_1), \dots, (e_{r-1}, h_{r-1})$ from table \mathcal{T} is the following operation. For every $0 \leq i \leq r-1$, remove every entry below e_i in h_{i+1} 's list in \mathcal{T} , i.e., move the bottom of h_{i+1} 's list up to e_i (from e_{i+1}). Then for each agent p who was just removed from h_{i+1} 's preference list, remove h_{i+1} from p 's list.

In Phase 2, exposed rotations are eliminated from the preference table one by one until some list becomes empty, which means that there is no stable matching, or no rotation is exposed in the table, which means that the list of every agent contains exactly one agent and we get a stable matching.

Proposition 4 (Irving 1994, Corollary 3.2). *If some agent obtains an empty preference list in Phase 2, then there is no stable matching.*

There are two types of rotations, which can be used to characterize all stable matchings. Let D denote the execution tree when Phase 2 is executed in all possible ways on the Phase 1 table. Note that each node x in D represents a table $\mathcal{T}(x)$, which is the current table state of the algorithm at node x . Let R be the set of all rotations which are exposed in some table of some node in D .

Definition 3 (Dual and singleton rotations). If $\rho \in R$ with $\rho = ((e_0, h_0), (e_1, h_1), \dots, (e_{r-1}, h_{r-1}))$, then the *negation* of ρ is defined as $\neg\rho = ((h_0, e_{r-1}), (h_1, e_0), \dots, (h_{r-1}, e_{r-2}))$.

If $\neg\rho \in R$, then we call ρ and $\neg\rho$ a *dual pair* of rotations. Any rotation without a dual is called a *singleton rotation*.

Definition 4. A subset $C \subseteq R$ of rotations is *complete* if it contains

- (i) all singleton rotations and
- (ii) exactly one rotation from each dual pair of rotation.

A binary relation \triangleright on R is defined as follows. For each pair $\{\rho_1, \rho_2\}$ of rotations, if the elimination of ρ_1 is necessary for ρ_2 to be exposed (i.e., the elimination of ρ_1 precedes the exposition of ρ_2 on every path in D), then we say that ρ_1 precedes ρ_2 , written as $\rho_1 \triangleright \rho_2$. The partial order \triangleright is the reflexive closure of the precedence relation \triangleright .

Note that the rotation poset (R, \triangleright) is an $O(n^2)$ -sized representation of all stable matchings. It follows from Gusfield (1988) that one can compute R in $O(n^3 \cdot \log(n))$ time.

It follows that in order to solve our problem, it suffices to search for a complete and closed subset of rotations for the second input profile \mathcal{P}_2 such that the corresponding stable matching is closest to M_1 .

Proposal sets and rotation weights

In this section, we study the influence of a rotation elimination on the distance between the ultimate resulting matching and the initial matching. In order to do so, we first define proposal sets for preference tables and weights of rotations. Here, the weight of a rotation shall capture how many pairs of the initial matching we can additionally obtain if we eliminate this rotation. Then we study the properties of the thus introduced rotation weights.

Imagine that in a preference table, every agent proposes to the first agent in their current preference list. Then, we get a set of ordered pairs, where the agent indicated by the first component of each pair proposes to the agent indicated by the second component in the pair.

Definition 5 (Proposal sets). For each preference table \mathcal{T} , the *proposal set* for \mathcal{T} is defined as $S_{\mathcal{T}} := \{(i, j) \mid i \in U \text{ and } j \text{ is the first agent on } i\text{'s list in } \mathcal{T}\}$, where (i, j) represents a *proposal pair* $i \rightarrow j$.

For each matching M , the *proposal set* S_M for M is defined as $S_M := \{(i, j), (j, i) \mid \{i, j\} \in M\}$.

By the definition of proposal sets, for each matching M we have

$$\begin{aligned} \text{dist}(M_1, M) &= |M_1 \Delta M| \\ &= |M_1| + |M| - 2|M_1 \cap M| \quad (1) \\ &= |M_1| + |M| - |S_{M_1} \cap S_M|. \end{aligned}$$

Hence, we are looking for a matching M_2 which is stable with respect to profile \mathcal{P}_2 such that

$$|S_{M_1} \cap S_{M_2}| \geq |M_1| + |M_2| - k. \quad (2)$$

Since every complete and closed rotation subset contains all singleton rotations and since all singleton rotations can be eliminated from the Phase 1 table before all dual rotations (Gusfield 1988), we can first eliminate all singleton rotations. By Proposition 4, if after eliminating all singleton rotations, some agent obtains an empty preference list, then we can immediately conclude that the given profile does not admit any stable matching. Thus, in the following, we mainly work with $R_2 \subseteq R$ which is the set of all dual rotations.

To measure the benefit of eliminating a rotation, we define the following.

Definition 6. For each rotation $\rho \in R$ with $\rho = (e_0, h_0), (e_1, h_1), \dots, (e_{r-1}, h_{r-1})$, the set of proposal pairs gained (lost) by eliminating rotation ρ is defined as follows.

$$S_{\rho}^{+} := \{(e_0, h_1), (e_1, h_2), \dots, (e_{r-1}, h_0)\},$$

$$S_{\rho}^{-} := \{(e_0, h_0), (e_1, h_1), \dots, (e_{r-1}, h_{r-1})\}.$$

Further, let $w^{+}(\rho) := |S_{\rho}^{+} \cap S_{M_1}|$ and $w^{-}(\rho) := |S_{\rho}^{-} \cap S_{M_1}|$ be the number of proposal pairs gained and lost by the elimination of rotation ρ .

Observe that the two sets S_{ρ}^{+} and S_{ρ}^{-} are independent of the table \mathcal{T} and all of the proposal pairs of agents not involved in the rotation ρ remain the same before and after the elimination of ρ .

Next, we prove that the weights of a dual pair of rotations are complementary.

Lemma 1. *Let $\rho \in R_2$, then $w^{+}(\rho) = w^{-}(\neg\rho)$.*

Reduction to WCS

In this subsection we give a reduction from INCREMENTAL STABLE ROOMMATES to WEIGHTED CONFLICT-FREE CLOSED SUBSET in which the conflict graph is a union of disjoint edges. In order to do so, we show that the distance between the target (i.e., initial) matching M_1 and a matching M_C resulting from the elimination of a complete and closed set C of rotations for \mathcal{P}_2 (if such set exists) depends only on $\sum_{\rho \in C} w^-(\rho)$.

In the remainder of the section, let S_0 be the proposal set for the table obtained from the Phase 1 table for \mathcal{P}_2 followed by elimination of all singleton rotations to it.

Before we continue with the procedure, we compare the sizes of M_C and M_1 ; recall that M_C is a matching resulting from eliminating a complete and closed set $C \subseteq R$. By Proposition 3, every agent that has non-empty list after Phase 1 must be matched under M_C and these agents are exactly those agents who hold some proposals in S_0 . Consequently, for each agent x that is matched under M_1 but does not hold a proposal under S_0 it holds that $\{x, M_1(x)\} \in M_1 \Delta M_C$. Thus, we define the following data reduction rule.

Reduction Rule 1. For each agent x matched under M_1 , if for each agent y it holds that $(x, y) \notin S_0$, then delete $\{x, M_1(x)\}$ from M_1 and decrease k by one.

Lemma 2. *Reduction Rule 1 is sound and can be executed in $O(n^2)$ time. Moreover, if Reduction Rule 1 does not apply, then for each stable matching M of profile \mathcal{P}_2 it holds that $|M_1| \leq |M| = |S_0|/2$.*

Next, we upper-bound the sum of weights of a complete and closed subset C which can result in a desired matching (i.e., $\text{dist}(M_1, M_C) \leq k$).

Lemma 3. *Let C be a complete and closed subset of rotations in R_2 and let M_C be the stable matching associated with C . Then, the following holds.*

- (i) $|S_{M_1} \cap S_{M_C}| = |S_{M_1} \cap S_0| + \sum_{\sigma \in R_2} w^+(\sigma) - 2 \cdot \sum_{\rho \in C} w^-(\rho)$.
- (ii) $\text{dist}(M_1, M_C) \leq k$ if and only if $\sum_{\rho \in C} w^-(\rho) \leq (|S_{M_1} \cap S_0| + \sum_{\sigma \in R_2} w^+(\sigma) - |M_1| - |M_C| + k)/2$.
- (iii) $|S_{M_1} \cap S_0| + \sum_{\sigma \in R_2} w^+(\sigma) - |M_1| - |M_C| \leq 0$.

This allows us to reduce the given instance (in polynomial time) to an instance of WCS as follows (also see Example 1).

Construction 1. *Finally, we arrive at the following instance of WEIGHTED CONFLICT-FREE CLOSED SUBSET:*

- (i) *Apply Reduction Rule 1 in $O(n^2)$ time, and compute in $O(n^3 \cdot \log n)$ time the rotation poset (R, \succeq) for \mathcal{P}_2 .*
- (ii) *Compute S_0 for the profile \mathcal{P}_2 .*
- (iii) *Let G be a graph on R_2 in which two elements of R_2 are adjacent if they form a dual pair of rotations (consequently, G is a union of $\ell = |R_2|/2$ disjoint edges).*
- (iv) *The weight function w is defined as w^- .*
- (v) *The budget b on the sum of weights is $b := (|S_{M_1} \cap S_0| + \sum_{\sigma \in R_2} w^+(\sigma) - |M_1| - |S_0|/2 + k)/2$.*



Figure 1: Left: A diagram for the partial order \succeq on R_2 . Right: The constructed graph G with weights next to the vertices.

Note that b is derived from Lemma 3(ii) such that we are searching for a complete and closed subset C of rotations whose sum of weights is upper-bounded by b . To see this, since $|S_M| = |S_0|/2$ (Lemma 2), the budget b is in fact equal to $(|S_{M_1} \cap S_0| + \sum_{\sigma \in R_2} w^+(\sigma) - |M_1| - |M| + k)/2$, where M is an arbitrary stable matching for \mathcal{P}_2 . Thus, by Lemma 3(ii), the budget b is at most $k/2$. Note also that the sum of the weights $w^-(\rho)$ is upper-bounded by n and thus the reduction presented above is a polynomial (many-to-one) reduction.

Example 1. Let $M_1 = \{\{1, 7\}, \{2, 3\}, \{4, 6\}, \{5, 8\}\}$ and $k = 4$. The profile \mathcal{P}_2 is from Gusfield (1988). The full analysis can be found in the full version of this paper (Bredereck et al. 2019). We start with the Phase 2 table where all singleton rotations (ρ_1 and ρ_3) are eliminated.

1	6	5	4	5	7	1	8	
2	3			6	8	4	5	1
3	2			7	4	5		
4	1	7		8	5	6		

Now the set of all dual rotations is $R_2 = \{\rho_2, \rho_4, \rho_5, \rho_6\}$, where $\rho_2 = \neg\rho_6$ and $\rho_4 = \neg\rho_5$. The partial order \succeq on R_2 and the constructed instance $((R = R_2, \succeq), G = (R_2, E), w, \ell = 2, b)$ of WCS are shown in Figure 1. The edge set E consists of two edges $\{\rho_2, \rho_6\}$ and $\{\rho_4, \rho_5\}$ since they are dual pairs of rotations. We take $\rho_4 = ((1, 6), (8, 5))$ as an example to show how to define the weight function w . Since $S_{\rho_4}^- = \{(1, 6), (8, 5)\}$, we have $w(\rho_4) = w^-(\rho_4) = |S_{\rho_4}^- \cap S_{M_1}| = |\{(8, 5)\}| = 1$. Similarly, we get $w(\rho_2) = w(\rho_5) = w(\rho_6) = 0$. The budget $b = (3 + 1 - 4 - 4 + 4)/2 = 0$. The task is to find a closed subset $C \subseteq R_2$ of size $\ell = 2$ which is independent in G such that $\sum_{c \in C} w(c) \leq 0$. It is easy to see that $C = \{\rho_2, \rho_5\}$ is the only solution. By eliminating ρ_2 and ρ_5 , we get matching $M_{\{\rho_1, \rho_2, \rho_3, \rho_5\}} = \{\{1, 6\}, \{2, 3\}, \{7, 4\}, \{8, 5\}\}$. It is easy to check that $\text{dist}(M_1, M_{\{\rho_1, \rho_2, \rho_3, \rho_5\}}) = 4 = k$.

Solving WCS when the conflict graph consists of ℓ disjoint cliques. Now, to prove Theorem 1, we only need to show the following; recall that the budget b , as defined in the construction, is at most $k/2$.

Lemma 4. *For graphs G with exactly ℓ cliques, WCS can be solved in $O((\Delta(G) + 1))^b \cdot |R|^2)$ time.*

We finally prove our main theorem.

Proof of Theorem 1. By Construction 1, in polynomial time, we construct an instance for WCS, where the budget b is upper-bounded by $k/2$ and the conflict graph consists of $|R_2|/2$ edges. Note that by Corollary 5.1 of Gusfield (1988)

we have that $|R_2| = O(n^2)$. By Lemma 4, we can solve this instance and thus our problem in $O(2^k \cdot n^4)$ time. \square

Hard Cases

Throughout this section, we are using the following non-standard “incremental” (resp. “decremental”) variants of the INDEPENDENT SET (resp. CLIQUE) problem to show parameterized intractability.

Our first problem asks for an independent set of size h for some graph in the case when an independent set of size h for the graph minus an edge is already known.

Our second problem asks for a size- h clique with pendant edges for some graph in the case when a size- h clique with pendant edges for the graph with an additional edge is known. A *clique with pendant edges for a graph G* is a subset $V' \subset V(G)$ of vertices such that V' forms a *clique* in G (each two vertices in V' are adjacent) and each vertex in V' has at least one neighbor outside V' . The size of a *clique with pendant edges* is defined as the number of vertices in the clique.

Lemma 5. EDGE-INCREMENTAL INDEPENDENT SET and EDGE-DECREMENTAL CLIQUE WITH PENDANT EDGES are NP-hard and $W[1]$ -hard with respect to h .

INCREMENTAL STABLE ROOMMATES without ties

To show that INCREMENTAL STABLE ROOMMATES without ties is NP-hard, we identify a relation of it to an egalitarian variant of stable matching where the egalitarian cost is minimized. Here, the *egalitarian cost* of a matching M is defined as the sum of the ranks of the agents with respect to their partners, and the *rank* of an agent x with respect to its partner $M(y)$ is equal to the number agents that x prefer over y . Feder (1992) showed that finding a stable matching with minimum egalitarian cost is NP-hard for STABLE ROOMMATES, even for complete preferences without ties (see Chen et al. (2018) for fixed-parameter tractability results on this problem).

The original hardness proof by Feder (1992) is to reduce from VERTEX COVER, which, given an undirected graph G and an integer $h' \in \mathbb{N}$, asks whether G admits a *vertex cover* of size h' , i.e., a size- h' vertex subset of $V' \subseteq V(G)$ such that each edge in $E(G)$ is incident to at least one vertex from V' . The basic idea behind the reduction is that putting a vertex to the solution set is equivalent to increasing the egalitarian cost by one. This correspondence can also be achieved in INCREMENTAL STABLE ROOMMATES by choosing an initial matching M_1 which is associated to an empty vertex set, thus finding a stable matching M_2 closest to M_1 is equivalent to finding a vertex cover of minimum size. Note that VERTEX COVER and INDEPENDENT SET are dual to each other, i.e., a vertex subset V' is a vertex cover of size h' if and only if $V(G) \setminus V'$ is an independent set of size $|V(G)| - h'$. Due to this and since INDEPENDENT SET is $W[1]$ -hard with respect to $|V| - h'$, we will directly reduce from INDEPENDENT SET, also showing parameterized intractability for our problem.

Proposition 5. INCREMENTAL STABLE ROOMMATES without ties is NP-hard and $W[1]$ -hard with respect to $k' = |M_1 \cap M_2|$.

INCREMENTAL STABLE MARRIAGE with ties

We show that INCREMENTAL STABLE MARRIAGE becomes intractable when ties are allowed even if the two preference profiles \mathcal{P}_1 and \mathcal{P}_2 are almost identical. The following result is achieved by a parameterized reduction from EDGE-DECREMENTAL CLIQUE WITH PENDANT EDGES.

Theorem 2. INCREMENTAL STABLE MARRIAGE with ties is $W[1]$ -hard with respect to k , even if $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$.

In the following, we show that with respect to the number of common pairs between the target stable matching and the initial stable matching the problem is parameterized intractable, even if the two input profiles differ by only two swaps. The corresponding parameterized reduction is from INDEPENDENT SET.

Theorem 3. INCREMENTAL STABLE MARRIAGE with ties is $W[1]$ -hard with respect to $k' = |M_1 \cap M_2|$ of common pairs, even if $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 2$.

Finally, we show that even a single swap in the preference list of one single agent makes the INCREMENTAL STABLE ROOMMATES problem $W[1]$ -hard with respect to k when ties are allowed. To show this result, we give a reduction from EDGE-INCREMENTAL INDEPENDENT SET. The construction idea is inspired by a reduction from VERTEX COVER to STABLE ROOMMATES with structured preferences (Bredereck et al. 2017), which, however, is relying on incomplete preferences and not showing parameterized intractability.

Theorem 4. INCREMENTAL STABLE ROOMMATES with ties and complete preferences is $W[1]$ -hard with respect to $k = |M_1 \Delta M_2|$, even if $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$. It remains NP-hard even if $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$, and the sought stable matching M_2 only needs to satisfy that $|M_1 \cap M_2| \geq 0$.

Conclusion

Motivated by dynamically changing preferences and the necessity to adapt the corresponding solutions, we introduced an “incremental view” on the computation of stable matchings. We believe that there are plenty of opportunities for future research, including, for instance, to study the role of parameters measuring the number of ties—in many hardness reductions this parameter is unbounded. We also left open whether INCREMENTAL STABLE ROOMMATES without ties is fixed-parameter tractable when parameterized by the swap distance between the two input preference profiles. Naturally, there are also future directions concerning more conceptual work, e.g., also studying further stability concepts in the context of our incremental model.

Acknowledgments

We thank anonymous reviewers of AAAI for their insightful comments. Main work was done while JC was with University of Warsaw, supported by ERC Horizon 2020 research and innovation programme (677651). JC was also supported by the WWTF research project (VRG18-012). Main work was done while DK was with TU Berlin, supported by the DFG project MaMu (NI 369/19). JL was supported by the DFG project AFFA (BR 5207/1 and NI 369/15).

References

- Abu-Khzam, F. N.; Egan, J.; Fellows, M. R.; Rosamond, F. A.; and Shaw, P. 2015. On the parameterized complexity of dynamic problems. *Theoretical Computer Science* 607:426–434.
- Aziz, H.; Biró, P.; Gaspers, S.; de Haan, R.; Mattei, N.; and Rastegari, B. 2016. Stable matching with uncertain linear preferences. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory*, 195–206.
- Bhattacharya, S.; Hofer, M.; Huang, C.; Kavitha, T.; and Wagner, L. 2015. Maintaining near-popular matchings. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, 504–515.
- Böckenhauer, H.; Burjons, E.; Raszyk, M.; and Rossmann, P. 2018. Reoptimization of parameterized problems. *CoRR* abs/1809.10578.
- Bredereck, R.; Chen, J.; Finnendahl, U. P.; and Niedermeier, R. 2017. Stable roommate with narcissistic, single-peaked, and single-crossing preferences. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory*, 315–330.
- Bredereck, R.; Chen, J.; Knop, D.; Luo, J.; and Niedermeier, R. 2019. Adapting stable matchings to evolving preferences. *CoRR* abs/1907.01375.
- Charikar, M.; Chekuri, C.; Feder, T.; and Motwani, R. 2004. Incremental clustering and dynamic information retrieval. *SIAM Journal on Computing* 33(6):1417–1440.
- Chen, J.; Hermelin, D.; Sorge, M.; and Yedidsion, H. 2018. How hard is it to satisfy (almost) all roommates? In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*, 35:1–35:15.
- Chen, J.; Niedermeier, R.; and Skowron, P. 2018. Stable marriage with multi-modal preferences. In *Proceedings of the 19th ACM Conference on Economics and Computation*, 269–286.
- Chen, J.; Skowron, P.; and Sorge, M. 2019. Matchings under preferences: Strength of stability and trade-offs. In *Proceedings of the 20th ACM Conference on Economics and Computation*, 41–59.
- Cseh, Á., and Manlove, D. F. 2016. Stable marriage and roommates problems with restricted edges: Complexity and approximability. *Discrete Optimization* 20:62–89.
- Drummond, J., and Boutilier, C. 2013. Elicitation and approximately stable matching with partial preferences. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 97–105.
- Feder, T. 1992. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences* 45(2):233–284.
- Genc, B.; Siala, M.; O’Sullivan, B.; and Simonin, G. 2017a. Finding robust solutions to stable marriage. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 631–637.
- Genc, B.; Siala, M.; Simonin, G.; and O’Sullivan, B. 2017b. Robust stable marriage. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 4925–4926.
- Genc, B.; Siala, M.; Simonin, G.; and O’Sullivan, B. 2019. Complexity study for the robust stable marriage problem. *Theoretical Computer Science* 775:76–92.
- Ghosal, P.; Kunysz, A.; and Paluch, K. E. 2017. The dynamics of rank-maximal and popular matchings. *CoRR* abs/1703.10594.
- Gusfield, D., and Irving, R. W. 1989. *The Stable Marriage Problem—Structure and Algorithms*. Foundations of Computing Series. MIT Press.
- Gusfield, D. 1988. The structure of the stable roommate problem: Efficient representation and enumeration of all stable assignments. *SIAM Journal on Computing* 17(4):742–769.
- Hartung, S., and Niedermeier, R. 2013. Incremental list coloring of graphs, parameterized by conservation. *Theoretical Computer Science* 494:86–98.
- Irving, R. W. 1985. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms* 6(4):577–595.
- Irving, R. W. 1994. Stable marriage and indifference. *Discrete Applied Mathematics* 48(3):261–272.
- Kanade, V.; Leonardos, N.; and Magniez, F. 2016. Stable matching with evolving preferences. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 60 of *LIPIcs*, 36:1–36:13.
- Krithika, R.; Sahu, A.; and Tale, P. 2018. Dynamic parameterized problems. *Algorithmica* 80(9):2637–2655.
- Luo, J.; Molter, H.; Nichterlein, A.; and Niedermeier, R. 2018. Parameterized dynamic cluster editing. In *Proceedings of the 38th International Conference on Foundations of Software Technology and Theoretical Computer Science*, 46:1–46:15.
- Mai, T., and Vazirani, V. V. 2018a. Finding stable matchings that are robust to errors in the input. In *Proceedings of the 26th Annual European Symposium on Algorithms*, 60:1–60:11.
- Mai, T., and Vazirani, V. V. 2018b. A generalization of Birkhoff’s theorem for distributive lattices, with applications to robust stable matchings. Technical report, arXiv:1804.05537 [cs.DM].
- Manlove, D. F. 2013. *Algorithmics of Matching Under Preferences*. WorldScientific.
- Marx, D., and Schlotter, I. 2010. Parameterized complexity and local search approaches for the stable marriage problem with ties. *Algorithmica* 58(1):170–187.
- Marx, D., and Schlotter, I. 2011. Stable assignment with couples: Parameterized complexity and local search. *Discrete Optimization* 8(1):25–40.
- Miyazaki, S., and Okamoto, K. 2017. Jointly stable matchings. In *Proceedings of the 28th International Symposium on Algorithms and Computation*, 56:1–56:12.
- Nimbhorkar, P., and Rameshwar, V. 2019. Dynamic rank-maximal and popular matchings. *Journal of Combinatorial Optimization* 37(2):523–545.
- Schieber, B.; Shachnai, H.; Tamir, G.; and Tamir, T. 2018. A theory and algorithms for combinatorial reoptimization. *Algorithmica* 80(2):576–607.