

Theory of and Experiments on Minimally Invasive Stability Preservation in Changing Two-Sided Matching Markets

Niclas Boehmer, Klaus Heeger, Rolf Niedermeier

Technische Universität Berlin, Faculty IV, Algorithmics and Computational Complexity, Berlin, Germany
 niclas.boehmer@tu-berlin.de, heeger@tu-berlin.de, rolf.niedermeier@tu-berlin.de

Abstract

Following up on purely theoretical work, we contribute further theoretical insights into adapting stable two-sided matchings to change. Moreover, we perform extensive empirical studies hinting at numerous practically useful properties. Our theoretical extensions include the study of new problems (that is, incremental variants of ALMOST STABLE MARRIAGE and HOSPITAL RESIDENTS), focusing on their (parameterized) computational complexity and the equivalence of various change types (thus simplifying algorithmic and complexity-theoretic studies for various natural change scenarios). Our experimental findings reveal, for instance, that allowing the new matching to be blocked by a few pairs significantly decreases the difference between the old and the new matching.

1 Introduction

In our dynamic world, change is omnipresent in society and business.¹ Typically, there is no permanent stability. We address this issue in the context of stable matchings in two-sided matching markets and their adaptivity to change. Consider as an example the dynamic nature of centrally assigning students to public schools. Here, students are matched to schools, trying to accommodate the students' preferences over the schools as well as possible. However, due to students reallocating or deciding to visit a private school, according to Feigenbaum et al. (2020), in New York typically around 10% of the students drop out after a first round of assignments, triggering some readjustments in the school-student matchings in a further round.

Matching students to schools can be modeled as an instance of the HOSPITAL RESIDENTS problem, where we are given a set of residents and hospitals each with preferences over the agents from the other set. One wants to find a "stable" assignment of each resident to at most one hospital such that a given capacity for each hospital is respected. To model the task of adjusting a matching to change, Bredereck et al. (2020) introduced the problem, given a stable matching with respect to some initial preference profile, to find a new

matching which is stable with respect to an updated preference profile (where some agents performed swaps in their preferences) and which is as similar as possible to the given matching. They referred to this as the "incremental" scenario and studied the computational complexity of this question for STABLE MARRIAGE and STABLE ROOMMATES (both being one-to-one matching problems).

In this work, we address multiple so-far unstudied aspects of our introductory school choice example. First, we theoretically and experimentally relate different types of changes to each other, including swapping two agents in some preference list (as studied by Bredereck et al. (2020)) and deleting an agent (as in our introductory example). Second, we initiate the study of the incremental variant of many-to-one stable matchings (HOSPITAL RESIDENTS). Third, as perfect stability might not always be essential, for instance, in large markets, we introduce the incremental variant of ALMOST STABLE MARRIAGE (where the new matching is allowed to be blocked by few agent pairs) and study its computational complexity and practical impact. Fourth, we experimentally analyze how many adjustments are typically needed when a certain amount of change occurs; moreover, we give some recommendations to market makers for adapting stable matchings to change.

1.1 Related Work

We are closest to the purely theoretical work of Bredereck et al. (2020), using their formulation of incremental stable matching problems (we refer to their related work section for an extensive discussion of related and motivating literature before 2020). Among others, they proved that INCREMENTAL STABLE MARRIAGE is polynomial-time solvable but is NP-hard (and W[1]-hard parameterized by the allowed change between the two matchings) if the preferences may contain ties. We complement and enhance some of Bredereck et al.'s findings, focusing on two-sided markets and contributing extensive experiments.

Besides the work of Bredereck et al. (2020), there are several other works dealing with adapting a (stable) matching to a changing agent set or changing preferences (Bhattacharya et al. 2015; Kanade, Leonardos, and Magniez 2016; Ghosal, Kunysz, and Paluch 2020; Nimbhorkar and Rameshwar 2019; Feigenbaum et al. 2020; Gajulapalli et al. 2020). Closest to our work, Gajulapalli et al. (2020) de-

signed polynomial-time algorithms for two variants of an incremental version of HOSPITAL RESIDENTS where the given matching is resident-optimal (unlike in our setting) and in the updated instance either new residents are added or the quotas of some hospitals are modified.

Sharing a common motivation with our work, there is a rich body of studies concerning dynamic matching markets mostly driven by economists (Damiano and Lam 2005; Akbarpour, Li, and Gharan 2020; Baccara, Lee, and Yariv 2020; Liu 2021). In the context of matching under preferences, a frequently studied exemplary (online) problem is that agents arrive over time and want to be matched as soon as possible in an—also in the long run—stable way (reassignments are not allowed) (Liu 2021; Doval 2021).

1.2 Our Contributions

On the theoretical side, while Brederick et al. (2020) focused on swapping adjacent agents in preference lists, we consider three further natural types of changes: the deletion and addition of agents and the complete replacement of an agent’s preference list. These different change types model different kinds of real-world scenarios; however, as one of our main theoretical results, we prove in Section 3 that all four change types are equivalent, thus allowing us to transfer both algorithmic and computational hardness results from one type to another.

Motivated by the polynomial-time algorithm of Brederick et al. (2020) for INCREMENTAL STABLE MARRIAGE, in Section 4 we study the related problem INCREMENTAL ALMOST STABLE MARRIAGE (where the new matching may admit few blocking pairs). We show that INCREMENTAL ALMOST STABLE MARRIAGE is NP-hard and establish parameterized tractability and intractability results. Moreover, motivated by the observation that, in practice, also many-to-one matching markets may change, we consider INCREMENTAL HOSPITAL RESIDENTS in Section 5. We show that the problem is polynomial-time solvable. However, if preferences may contain ties, then it becomes NP-hard and W[1]-hard when parameterized by the number of hospitals; still, we can identify several (fixed-parameter) tractable cases.

On the experimental side (Section 6), we perform an extensive study, among others taking into account the four different change types discussed above. For instance, we investigate the relation between the number of changes and the symmetric difference between the old and new stable matching. We observe that often already very few random changes require a major restructuring of the matching. One way to circumvent this problem is to allow that the new matching might be blocked by a few agent pairs. Moreover, reflecting its popularity, we compute the input matching using the Gale-Shapley algorithm (Gale and Shapley 1962) and observe that, in this case, computing the output matching also with the Gale-Shapley algorithm produces a close to optimal solution.

The proofs (or their completions) for results marked by (★) and major parts of our experimental study can be found in a full version of our work (Boehmer, Heeger, and Niedermeier 2021).

2 Preliminaries

An instance of the STABLE MARRIAGE WITH TIES (SM-T) problem consists of two sets U and W of agents and a preference profile \mathcal{P} containing a preference relation for each agent. Following conventions, we refer to the agents from U as men and to the agents from W as women. We denote the set of all agents by $A := U \cup W$. Each man $m \in U$ accepts a subset $\text{Ac}(m) \subseteq W$ of women, and each woman w accepts a subset $\text{Ac}(w) \subseteq U$ of men. The preference relation \succsim_a of agent $a \in A$ is a weak order of the agents $\text{Ac}(a)$ that agent a accepts. For two agents $a', a'' \in \text{Ac}(a)$, agent a *weakly prefers* a' to a'' if $a' \succsim_a a''$. If a both weakly prefers a' to a'' and a'' to a' , then a is *indifferent* between a' and a'' and we write $a' \sim_a a''$. If a weakly prefers a' to a'' but does not weakly prefer a'' to a' , then a *strictly prefers* a' to a'' and we write $a' \succ_a a''$. If the preference relation of an agent a is a strict order, that is, there are no two agents such that a is indifferent between the two, then we say that a has *strict preferences* and denote a ’s preference relation as \succ_a . In this case, we use the terms “strictly prefer” and “prefer” interchangeably. STABLE MARRIAGE (SM) is the special case of SM-T where all agents have strict preferences. For two preference relations \succsim and \succsim' , the *swap distance* between \succsim and \succsim' is the number of agent pairs that are ordered differently by the two relations, i.e., $|\{\{a, b\} : a \succ b \wedge b \not\succ' a\}| + |\{\{a, b\} : a \sim b \wedge \neg a \sim' b\}|$; if both relations are defined on different sets, then we define the swap distance to be infinity. For two strict preference relations \succ and \succ' , the swap distance yields the minimum number of swaps of adjacent agents needed to transform \succ into \succ' . For two preference profiles \mathcal{P}_1 and \mathcal{P}_2 on the same set of agents, $|\mathcal{P}_1 \oplus \mathcal{P}_2|$ denotes the summed swap distance between the two preference relations of each agent.

A *matching* M is a set of pairs $\{m, w\}$ with $m \in \text{Ac}(w)$ and $w \in \text{Ac}(m)$ where each agent appears in at most one pair. For two matchings M and M' , the *symmetric difference* is $M \Delta M' = (M \setminus M') \cup (M' \setminus M)$. In a matching M , an agent a is *matched* if a appears in one pair, i.e., $\{a, a'\} \in M$ for some $a' \in A \setminus \{a\}$; otherwise, a is *unmatched*. A matching is *perfect* if each agent is matched. For a matching M and a matched agent $a \in A$, we denote by $M(a)$ the partner of a in M , i.e., $M(a) = a'$ if $\{a, a'\} \in M$. For an unmatched agent $a \in A$, we set $M(a) := \emptyset$. All agents $a \in A$ strictly prefer any agent from $\text{Ac}(a)$ to being unmatched (thus, we have $a' \succ_a \emptyset$ for $a' \in \text{Ac}(a)$).

A pair $\{u, w\}$ with $u \in U$ and $w \in W$ *blocks* a matching M if m and w accept each other and strictly prefer each other to their partners in M , i.e., $m \in \text{Ac}(w)$, $w \in \text{Ac}(m)$, $m \succ_w M(w)$, and $w \succ_m M(m)$. A matching M is *stable* if it is not blocked by any pair. SM and SM-T ask whether there exists a stable matching of the agents A with respect to preference profile \mathcal{P} .

We also consider a generalization of SM called ALMOST STABLE MARRIAGE (ASM), where as an additional part of the input we are given an integer b and the question is whether there is a matching admitting at most b blocking pairs. Furthermore, we study the HOSPITAL RESIDENTS (HR) problem, a generalization of SM where we are given a set R of residents and a set H of hospitals and

agents from both sets have preferences over a set of acceptable agents from the other set and each hospital $h \in H$ has an upper quota $u(h)$. In a matching, each resident can appear in at most one pair, while each hospital h can appear in at most $u(h)$ pairs. In this context, we slightly adapt the definition of a blocking pair and say that a resident-hospital pair $\{r, h\}$ blocks a matching M if both r and h accept each other, r prefers h to $M(r)$, and h is matched to less than $u(h)$ residents in M or prefers r to one of the residents matched to it.

Our work focuses on “incrementalized versions” of the discussed two-sided stable matching problems. For SM/(SM-T), this reads as follows:

INCREMENTAL STABLE MARRIAGE [WITH TIES]
(ISM/[ISM-T])

Input: A set $A = U \cup W$ of agents, two preference profiles \mathcal{P}_1 and \mathcal{P}_2 containing the strict [weak] preferences of all agents, a stable matching M_1 in \mathcal{P}_1 , and an integer k .

Question: Is there a matching M_2 that is stable in \mathcal{P}_2 such that at most k edges appear in only one of M_1 and M_2 , i.e., $|M_1 \Delta M_2| \leq k$?

IHR and IHR-T are defined analogously. IASM [IASM-T] is defined as ISM [ISM-T] with the difference that we are given an additional integer b as part of the input and the question is whether there is a matching M_2 that admits at most b blocking pairs in \mathcal{P}_2 such that $|M_1 \Delta M_2| \leq k$.

3 Equivalence of Different Types of Changes

Bredereck et al. (2020) focused on the case where the preference profile \mathcal{P}_2 arises from \mathcal{P}_1 by performing some swaps in the preferences of some agents (we refer to this as **Swap**). However, there are many more types of changes: Allowing for more radical changes, denoted by **Replace**, we count the number of agents whose preferences changed (here in contrast to **Swap**, we also allow that the set of acceptable partners may change). Next, recall that in our introductory example from school choice children leave the matching market, which corresponds to agents getting deleted. We denote this type of change by **Delete**—formally, we model the deletion of an agent by setting its set of acceptable partners in \mathcal{P}_2 to \emptyset . Moreover, children leaving one market might enter a new one, which corresponds to agents getting added (**Add**). Formally, we model the addition of an agent a by already including it in \mathcal{P}_1 , but with $\text{Ac}(a) = \emptyset$ in \mathcal{P}_1 . The goal of this section is to show that these four natural possibilities of how \mathcal{P}_2 may arise from \mathcal{P}_1 actually result in equivalent computational problems. More formally, we say that a type of change $\mathcal{X} \in \{\text{Delete, Add, Swap, Replace}\}$ linearly reduces to a type of change $\mathcal{Y} \in \{\text{Delete, Add, Swap, Replace}\}$ if any instance $\mathcal{I} = (A, \mathcal{P}_1, \mathcal{P}_2, M_1, k)$ of ISM(-T) where \mathcal{P}_1 and \mathcal{P}_2 differ by x changes of type \mathcal{X} can be transformed in linear time to an equivalent instance $\mathcal{I}' = (A', \mathcal{P}'_1, \mathcal{P}'_2, M'_1, k')$ of ISM(-T) with \mathcal{P}'_1 and \mathcal{P}'_2 differing by $\mathcal{O}(x)$ changes of type \mathcal{Y} . We call two change types \mathcal{X} and \mathcal{Y} linearly equivalent if both \mathcal{X} linearly reduces to \mathcal{Y} and \mathcal{Y} linearly reduces to \mathcal{X} .

Theorem 1 (★). *Swap, Replace, Delete, and Add are linearly equivalent for ISM and ISM-T.*

We only exemplarily show here how **Replace** can be linearly reduced to **Add**:

Lemma 1 (★). *Replace can be linearly reduced to Add.*

Proof. Let $\mathcal{I} = (A = U \cup W, \mathcal{P}_1, \mathcal{P}_2, M_1, k)$ be an instance of ISM(-T) for **Replace**. From this, we construct an instance $\mathcal{I}' = (A' = U' \cup W', \mathcal{P}'_1, \mathcal{P}'_2, M'_1, k')$ of ISM(-T) for **Add** as follows. Let A_{repl} be the set of agents with different preferences in \mathcal{P}_1 and \mathcal{P}_2 , and let $A_{\text{repl}}^* := A_{\text{repl}} \cup \{M_1(a) : a \in A_{\text{repl}} \wedge M_1(a) \neq \emptyset\}$ be the set of these agents and their partners in M_1 . To construct \mathcal{I}' , we start by adding all agents from A to A' and set the preferences of all agents in \mathcal{P}'_1 and \mathcal{P}'_2 to be their preferences in \mathcal{P}_1 (the preferences of some of these agents will be modified slightly in the following). Moreover, for each $a \in A_{\text{repl}}^*$, we add to A' a “binding” agent b_a and a “clone” c_a . Agent c_a has empty preferences in \mathcal{P}'_1 and has a ’s preferences from \mathcal{P}_2 in \mathcal{P}'_2 . We modify the preferences of all so far added agents such that c_a appears directly before a (or is tied with a if we have an instance with ties). Agent b_a has empty preferences in \mathcal{P}'_1 , only finds a acceptable in \mathcal{P}'_2 , and we modify the preferences of a in both \mathcal{P}'_1 and \mathcal{P}'_2 such that a prefers b_a to all other agents.

The idea behind the construction is as follows. We add b_a in \mathcal{P}'_2 which forces M'_2 to contain $\{a, b_a\}$ and further add agent c_a , who “replaces” a in \mathcal{P}'_2 and has a ’s changed preferences. However, this construction does not directly work: Let $m \in A_{\text{repl}}^* \cap U$ and $w = M_1(m)$. Unfortunately, adding the edge $\{m, w\}$ to M_2 corresponds to adding the edge $\{c_m, c_w\}$ to M'_2 , which leads to an increase of $|M'_1 \Delta M'_2|$ but not of $|M_1 \Delta M_2|$. In order to cope with this, we replace the edge $\{c_m, c_w\}$ by an *edge gadget* consisting of multiple agents: For each man $m \in A_{\text{repl}}^* \cap U$ matched by M_1 to a woman w , we introduce agents as depicted in Figure 1 and modify the preferences of c_m and c_w by replacing w and m by a_m^{lm} and a_w^{rm} , respectively.² The newly introduced agents from this gadget have empty preferences in \mathcal{P}'_1 and preferences as depicted in Figure 1 in \mathcal{P}'_2 except for agents a_m^{rm} and a_w^{lm} who have their depicted preferences in both \mathcal{P}'_1 and \mathcal{P}'_2 . We set $M'_1 := M_1 \cup \{\{a_m^{\text{rm}}, a_w^{\text{lm}}\} : \{m, w\} \in M_1 \wedge m \in A_{\text{repl}}^* \cap U\}$ and $k' := k + |A_{\text{repl}}^*| + 7k^*$ with $k^* := |\{\{m, w\} \in M_1 : m, w \in A_{\text{repl}}^*\}|$.

Next, we show the correctness of the forward direction of our reduction. Given a stable matching M_2 in \mathcal{P}_2 , we construct a stable matching M'_2 in \mathcal{P}'_2 with $|M'_1 \Delta M'_2| = |M_1 \Delta M_2| + |A_{\text{repl}}^*| + 7k^*$ as follows. We start with $M'_2 := M_1$. We first implement the adjustments corresponding to edges from $M_1 \Delta M_2$: Let $\{m, w\} \in M_2 \setminus M_1$. We delete the edges containing m and w from M'_2 (if there are any). Moreover, if $m, w \notin A_{\text{repl}}^*$, then we add $\{m, w\}$ to M'_2 . If $m \in A_{\text{repl}}^*$ and $w \notin A_{\text{repl}}^*$, then we add $\{c_m, w\}$. If $w \in A_{\text{repl}}^*$ and $m \notin A_{\text{repl}}^*$, then we add $\{m, c_w\}$. If $m, w \in A_{\text{repl}}^*$, then we add $\{c_m, c_w\}$. After these adjustments, it holds that $|M'_2 \Delta M'_1| = |M_2 \Delta M_1|$.

²We remark that this gadget is a concatenation of two parallel-edges gadgets used by Ceclárová and Fleiner (2005).

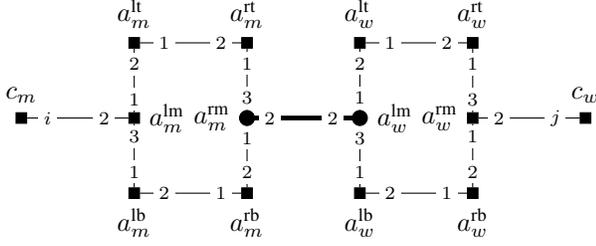


Figure 1: The edge gadget for edge $e = \{c_m, c_w\}$, where $m \in U$, $w \in W$, and m ranks w at the i -th rank, and w ranks m at the j -th rank. Squared agents have empty preferences in \mathcal{P}'_1 . The edge contained in M'_1 is bold. The numbers on the edges indicate the preferences of the agents: The number x closer to an agent a means that a ranks the other endpoint a' of the edge at rank x , i.e., there are $x - 1$ agents which a prefers to a' .

We now turn to the matching in the edge gadgets. For every edge $\{m, w\} \in M_1 \cap M_2$ with $m, w \in A_{\text{repl}}^*$, we delete $\{m, w\}$ from M'_2 and add edges $\{c_m, a_m^{\text{lm}}\}$, $\{a_m^{\text{lt}}, a_m^{\text{rt}}\}$, $\{a_m^{\text{lb}}, a_m^{\text{rb}}\}$, $\{a_w^{\text{lt}}, a_w^{\text{rt}}\}$, $\{a_w^{\text{lb}}, a_w^{\text{rb}}\}$, and $\{a_w^{\text{rm}}, c_w\}$. This contributes seven edges to $M'_1 \Delta M'_2$. For every edge $\{m, w\} \in M_1 \setminus M_2$ with $m, w \in A_{\text{repl}}^*$, we first delete edge $\{a_m^{\text{rm}}, a_w^{\text{lm}}\}$ from M'_2 . Subsequently, we make a case distinction based on whether m strictly prefers w to $M_2(m)$. If yes, then the stability of M_2 implies that w does not strictly prefer m to $M_2(w)$. Thus, we can add the edges $\{a_m^{\text{lb}}, a_m^{\text{rb}}\}$, $\{a_m^{\text{rm}}, a_m^{\text{rt}}\}$, $\{a_m^{\text{lt}}, a_m^{\text{lm}}\}$, $\{a_w^{\text{lb}}, a_w^{\text{rb}}\}$, $\{a_w^{\text{rm}}, a_w^{\text{rt}}\}$, and $\{a_w^{\text{lm}}, a_w^{\text{lw}}\}$, and the resulting matching is not blocked by $\{c_m, a_m^{\text{lm}}\}$. Otherwise, m does not strictly prefer w to $M_2(w)$. Thus, we can add the edges $\{a_m^{\text{lb}}, a_m^{\text{lm}}\}$, $\{a_m^{\text{rm}}, a_m^{\text{rb}}\}$, $\{a_m^{\text{lt}}, a_m^{\text{rt}}\}$, $\{a_w^{\text{lb}}, a_w^{\text{lm}}\}$, $\{a_w^{\text{rm}}, a_w^{\text{rb}}\}$, and $\{a_w^{\text{lt}}, a_w^{\text{rt}}\}$, and the resulting matching is not blocked by $\{a_w^{\text{rm}}, c_w\}$. This contributes seven edges to $M'_1 \Delta M'_2$. Thus, as we have k^* edge gadgets each contributing seven edges, we have $|M'_1 \Delta M'_2| = |M_1 \Delta M_2| + 7k^*$.

Lastly, for every $a \in A_{\text{repl}}^*$, we add the edge $\{a, b_a\}$ to M'_2 , which contributes $|A_{\text{repl}}^*|$ edges to $|M'_1 \Delta M'_2|$ leading to an overall symmetric difference of $|M'_1 \Delta M'_2| = |M_1 \Delta M_2| + |A_{\text{repl}}^*| + 7k^*$. It is easy to verify that M'_2 is stable in \mathcal{P}'_2 .

The reverse direction can be found in the full version. \square

Theorem 1 allows us to transfer algorithmic and hardness results for one type of change to another type. For example, the polynomial-time algorithm of Bredereck et al. (2020) for ISM for Swap implies that ISM can also be solved in polynomial time for Add, Delete, and Replace. Using similar constructions, it is also possible to prove that the different types of changes are equivalent for IHR (although, here, to model x changes of type \mathcal{X} more than $\mathcal{O}(x)$ changes of type \mathcal{Y} may be needed; e.g., in the above reduction from Replace to Add, modeling the replacement of a hospital h would need $u(h)$ binding residents b_h) and STABLE ROOMMATES (which is a generalization of SM where agents are not partitioned into men and women). However, Theorem 1

does not directly transfer to IASM; for instance, in the reduction from Lemma 1, M'_2 might “ignore” the added edge gadgets by allowing few of the edges to block M'_2 .

4 Almost Stable Marriage

Sometimes, it may be acceptable that “few” agent pairs block an implemented matching (for instance, in very large markets where agents might not even be aware that they are part of a blocking pair). In Section 6, we experimentally show that allowing that M_2 may be blocked by few agent pairs significantly decreases the number of necessary adjustments. We now show that, in contrast to ISM (Bredereck et al. 2020), IASM is computationally intractable:

Theorem 2 (★). *IASM is NP-hard and W[1]-hard when parameterized by $k + b + |\mathcal{P}_1 \oplus \mathcal{P}_2|$.*

To show Theorem 2, we devise a polynomial-time many-one reduction from LOCAL SEARCH ASM. In LOCAL SEARCH ASM, we are given an SM instance (U, W, \mathcal{P}) , a stable matching N in \mathcal{P} , and integers q , t , and z , and the question is whether there is a matching N^* of size at least $|N| + t$ admitting at most z blocking pairs such that $|N \Delta N^*| \leq q$. Gupta et al. (2020, Theorem 3) proved that LOCAL SEARCH ASM is NP-hard and W[1]-hard with respect to the combined parameter $q + t + z$, even if N^* needs to be a perfect matching. The general idea of the reduction behind Theorem 2 is to construct \mathcal{P}_1 from a LOCAL SEARCH ASM instance by adding a penalizing component and a set of catch men, and defining M_1 to be a stable matching containing N where each woman who is unmatched in N is matched to a catch man. Then, \mathcal{P}_2 differs from \mathcal{P}_1 in the preferences of all women unmatched in N who now prefer agents from the penalizing component to the catch men. After this change, these women need to be matched by M_2 (otherwise there will be too many blocking pairs with agents from the penalizing component) to agents from $U \cup W$ (otherwise the symmetric difference will exceed the budget). The details of the construction require some care, as allowing for some blocking pairs makes it more challenging to enforce how certain agents are matched in M_2 .

On the positive side, we provide XP-algorithms for all three single parameters:

Proposition 1 (★). *IASM is in XP when parameterized by any of k or b or $|\mathcal{P}_1 \oplus \mathcal{P}_2|$.*

Proof sketch. For parameter k , we guess the edges contained in $M_1 \Delta M_2$. For parameter b , we guess which agents form blocking pairs in M_2 and delete the mutual acceptability of these pairs in \mathcal{P}_2 . As the matching M_2 needs to be stable in this modified \mathcal{P}_2 , one can reduce finding the rest of the matching to the polynomial-time solvable WEIGHTED STABLE MARRIAGE problem (Feder 1992), which asks for a stable matching that maximizes some given edge weights. For parameter $|\mathcal{P}_1 \oplus \mathcal{P}_2|$, note that each swap can only create a single blocking pair for M_1 . Thus, if $|\mathcal{P}_1 \oplus \mathcal{P}_2| \leq b$, then we can simply set M_2 to M_1 . Otherwise, we use the XP-algorithm for the number b of blocking pairs. \square

We finally remark that while the XP-algorithm for the parameter k also works for IASM-T, IASM-T is NP-hard

even for $b = 0$ and $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$ (as Brederbeck et al. (2020) proved that ISM-T is NP-hard for $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$).

5 Incremental Hospital Residents

We start our study of the incremental variant of HOSPITAL RESIDENTS by observing that one can reduce IHR to the polynomial-time solvable WEIGHTED STABLE MARRIAGE problem (Feder 1992); this yields:

Proposition 2 (★). *IHR is solvable in $\mathcal{O}(n^{2.5} \cdot m^{1.5})$ time, where n is the number of residents and m is the number of hospitals.*

In the rest of this section, we focus on IHR-T. As IHR-T generalizes ISM-T, the results of Brederbeck et al. (2020) imply that IHR-T is NP-hard and W[1]-hard parameterized by k even for $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$. Thus, we focus on the parameters number n of residents and number m of hospitals.

For the number n of residents, we can bound the number of “relevant” hospitals by $\mathcal{O}(n^2)$. Subsequently guessing for each resident the hospital it is matched to yields:

Proposition 3 (★). *IHR-T is solvable in $\mathcal{O}(n^{2n} \cdot nm)$ time.*

Proposition 3 means fixed-parameter tractability with respect to n . In contrast to this, the number of hospitals is (presumably) not sufficient to gain fixed-parameter tractability, even if the two preference profiles differ only in one swap:

Theorem 3 (★). *Parameterized by the number m of hospitals, IHR-T is W[1]-hard even if $|\mathcal{P}_1 \oplus \mathcal{P}_2| = 1$.*

Proof (Construction). We reduce from the COM HR-T problem: Given an instance of HR-T, decide whether there is a stable matching which matches all residents. Boehmer and Heeger (2021, Proposition 8) showed that COM HR-T is W[1]-hard when parameterized by the number m of hospitals. Given an instance $\mathcal{I} = (R = \{r_1, \dots, r_n\} \cup H = \{h_1, \dots, h_m\}, \mathcal{P})$ of COM HR-T, let N be an arbitrary stable matching in \mathcal{I} (we assume that N does not match all residents, as we otherwise know that \mathcal{I} is a YES-instance).

To construct an instance of IHR-T, we first add $R \cup H$ to the set of agents. Subsequently, we add a penalizing component consisting of two hospitals h_1^* and h_2^* , both with upper quota one, and two hospitals \tilde{h}_1 and \tilde{h}_2 both with upper quota $n + 1$. We additionally add a resident r^* and two sets of $n + 1$ residents $\tilde{r}_1, \dots, \tilde{r}_{n+1}$ and $\tilde{r}'_1, \dots, \tilde{r}'_{n+1}$.

Turning to the agents’ preferences in \mathcal{P}_1 , all agents from $R \cup H$ have their preferences from \mathcal{P} , except that, for each resident, h^* is added at the end of her preferences. The preferences of the agents from the penalizing component are:

$$\begin{aligned} h_1^* &: r_1 \succ \dots \succ r_n \succ r^*; & r^* &: h_1^* \succ h_2^* \succ \tilde{h}_1; & h_2^* &: r^*; \\ \tilde{h}_1 &: r^* \succ \tilde{r}'_1 \succ \dots \succ \tilde{r}'_{n+1} \succ \tilde{r}_1 \succ \dots \succ \tilde{r}_{n+1}; \\ \tilde{h}_2 &: \tilde{r}_1 \succ \dots \succ \tilde{r}_{n+1} \succ \tilde{r}'_1 \succ \dots \succ \tilde{r}'_{n+1}; \\ \tilde{r}_i &: \tilde{h}_1 \succ \tilde{h}_2; & \tilde{r}'_i &: \tilde{h}_2 \succ \tilde{h}_1, & & i \in [n + 1]. \end{aligned}$$

Profile \mathcal{P}_2 equals \mathcal{P}_1 except that we swap h_2^* and \tilde{h}_1 in the preferences of r^* . Let i^* be the smallest index of a resident unmatched in N . We set $k := 2(n + 1)$ and

$$M_1 := N \cup \{\{r_{i^*}, h_1^*\}, \{r^*, h_2^*\}\}$$

$$\cup \{\{\tilde{r}_i, \tilde{h}_1\}, \{\tilde{r}'_i, \tilde{h}_2\} \mid i \in [n + 1]\}.$$

The correctness of the reduction crucially relies on the observation that M_2 needs to contain the edge $\{r^*, h_1^*\}$: Otherwise r^* is to be matched to \tilde{h}_1 , implying that all residents from the penalizing component need to be matched differently in M_2 than in M_1 , yielding $|M_1 \Delta M_2| > k$. From $\{r^*, h_1^*\} \in M_2$ it follows that all residents r_1, \dots, r_n are matched to hospitals from H in M_2 . Thus, M_2 induces a matching for the COM HR-T instance \mathcal{I} which matches all residents. \square

We leave open whether the (above shown) W[1]-hardness of IHR-T upholds for the parameter $m + k + |\mathcal{P}_1 \oplus \mathcal{P}_2|$.

On the positive side, devising an Integer Linear Program whose number of variables is upper-bounded in a function of m and some guessing as preprocessing, IHR-T admits an XP-algorithm for the number m of hospitals:

Proposition 4 (★). *IHR-T is in XP when parameterized by the number m of hospitals.*

In an SM-T instance, we say that two agents are of the same agent type if they have the same preference relation and all other agents are indifferent between them. One can interpret a hospital in an instance of HR-T as $u(h)$ agents of the same type and thus an HR-T instance as an instance of SM-T where agents from one side are of only m different agent types. This interpretation raises the question what happens when we parameterize ISM-T by the total number of agent types on both sides (and not only by the number of agent types on one of the sides as done in Theorem 3 and Proposition 4). We show that, in fact, this is enough to establish fixed-parameter tractability:

Proposition 5 (★). *ISM-T is solvable in $\mathcal{O}(2^{(t_U+1) \cdot (t_W+1)} \cdot n^{2.5})$ time, where t_U respectively t_W is the number of agent types of men respectively women in \mathcal{P}_2 .*

Proof sketch. Let T_U resp. T_W be the sets of agent types of men resp. women in \mathcal{P}_2 . We modify the instance by adding a new dummy men (women) type consisting of n men (women) who are indifferent among all women (men) and are ranked last by all women (men). We then iterate over all bipartite graphs G on $T_U \cup T_W$. We say that a matching M is compatible with G if M matches agents of type $\alpha \in T_U$ to agents of type $\beta \in T_W$ only if $\{\alpha, \beta\} \in E(G)$. We reject G if a matching M compatible with G can be unstable. To be precise, we reject G if there are two types $\alpha \in T_U$ and $\beta \in T_W$ such that there is an edge between α and some $\beta' \in T_W$ and an edge between β and some $\alpha' \in T_U$ such that agents of type α prefer agents of type β to agents of type β' and agents of type β prefer agents of type α to agents of type α' . If G is not rejected, then we construct a graph G^* on A from it by connecting agents of type $\alpha \in T_U$ and type $\beta \in T_W$ if and only if $\{\alpha, \beta\} \in E(G)$. Moreover, we assign all edges in G^* that appear in M_1 weight 1, and all other edges weight 0 if they contain a dummy agent and weight -1 otherwise. We compute a maximum weight matching M in G^* in $\mathcal{O}(n^{2.5})$ time (Duan and Su 2012) and return YES if M has weight at least $|M_1| - k$, and otherwise continue with the next graph G . \square

Notably, the above algorithm with minor modifications also works for the incremental variant of STABLE ROOMMATES WITH TIES.

6 Experiments

In this section, we consider different practical aspects.³ To keep the setup of our experiments simple, we focus on ISM, our most basic model. In Section 6.1, we analyze the relationship between the difference between \mathcal{P}_1 and \mathcal{P}_2 and the size $|M_1 \Delta M_2|$ of the symmetric difference between M_1 and M_2 . In Section 6.2, we study the trade-off between allowing M_2 to be blocked by some pairs and $|M_1 \Delta M_2|$. Lastly, in Section 6.3 we summarize further findings.

After having analyzed the theoretical relationship between different types of changes in Section 3, in this section, we compare the impact of the following three different types of changes (we always assume that all agents have complete and strict preferences):

Reorder A Reorder operation consists of permuting the preference list of an agent uniformly at random.

Delete A Delete operation consists of deleting an agent from the instance.

Swap A swap consists of swapping two adjacent agents in the preference relation of an agent. As sampling preference profiles that are at a certain swap distance from a given one is practically infeasible with more than 30 agents (Boehmer et al. 2021), we always perform the same number of swaps in the preferences of each agent: If we are to perform i Swap operations, then for each agent separately we replace its preferences by uniformly at random sampled preferences that are at swap distance i from its original preferences (using the procedure described by Boehmer et al. (2021)).

6.1 Stable Marriage

In this section, we analyze the relationship between the number of changes that are applied to \mathcal{P}_1 to obtain \mathcal{P}_2 and the size of the symmetric difference between the given matching M_1 that is stable in \mathcal{P}_1 and a stable matching M_2 in \mathcal{P}_2 .

Experimental Setup. For each of the three considered types of changes, for $r \in \{0, 0.01, 0.02, \dots, 0.3\}$ we sampled 200 STABLE MARRIAGE instances with 50 men and 50 women with random preferences (collected in the preference profile \mathcal{P}_1). For each of these instances, we set M_1 to be the men-optimal matching.⁴ Afterwards, we applied a uniformly at random sampled r -fraction of all possible changes of the considered type to profile \mathcal{P}_1 to obtain profile \mathcal{P}_2 . Subsequently, we computed a stable matching M_2 in \mathcal{P}_2 with minimum/maximum normalized symmetric difference $\frac{|M_1 \Delta M_2|}{|M_1| + |M_2|}$ to M_1 .⁵ We denote the solution with minimum symmetric difference as “Best” and the solution with

³For a detailed discussion of all mentioned experiments, please see the full version (Boehmer, Heeger, and Niedermeier 2021).

⁴We used the implementation of the Gale-Shapley algorithm of Wilde, Knight, and Gillard (2020).

⁵To compute this, we solved an (Integer) Linear Programming formulation of this problem as described in the full version

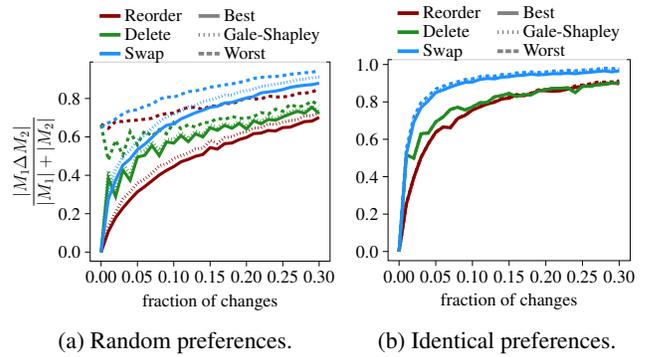


Figure 2: For different types of changes and ways to compute M_2 , average normalized symmetric difference between M_1 and M_2 for a varying fraction of change between \mathcal{P}_1 and \mathcal{P}_2 .

maximum symmetric difference as “Worst”. Moreover, we computed the men-optimal matching in \mathcal{P}_2 using the Gale-Shapley algorithm and denote this as “Gale-Shapley”. The results of this experiment are depicted in Figure 2a.

Evaluation. We start by focusing on the optimal solution (“Best”; solid line in Figure 2a). What stands out from Figure 2a is that already very few or even one change in \mathcal{P}_1 requires a fundamental restructuring of the given matching M_1 . To be precise, for Reorder, one reordering (which corresponds to a 0.01 fraction of changes) results in an average normalized symmetric difference between M_1 and M_2 of 0.1. For Swap, a 0.01-fraction of all swaps, which corresponds to making twelve random swaps per preference order (the total number of swaps is $\frac{n \cdot (n-1)}{2}$), results in an average normalized symmetric difference of 0.28, whereas a single swap per preference order already results in an average normalized symmetric difference of 0.05. For Delete, the effect was strongest, as deleting a single agent leads to an average normalized symmetric difference of 0.38. For Delete, Ashlagi, Kanoria, and Leshno (2017), Cai and Thomas (2021), Knuth, Motwani, and Pittel (1990) and Pittel (1989) offer some theoretical intuition of this phenomenon: Assuming that agents have random preferences (as in our experiments), with high probability in a men-optimal matching the average rank that a man has for the woman matched to him is $\log(n)$ (Knuth, Motwani, and Pittel 1990; Pittel 1989), whereas in an instance with n men and $n - 1$ women the average rank a man has for the woman matched to him is $\frac{n}{3 \log(n)}$ (Ashlagi, Kanoria, and Leshno 2017; Cai and Thomas 2021). Thus, if we delete a single woman from the instance (which happens with 50% probability when we delete a single agent), then already only to realize these average ranks, the given matching needs to be fundamentally restructured. Notably, if we delete two agents from the instance, which results only with a 25% probability in a higher number of men than women, then the minimum normalized symmetric difference between M_1 and M_2 is only 0.28.

(Boehmer, Heeger, and Niedermeier 2021) using Gurobi Optimization, LLC (2021).

While ISM is solvable in polynomial time, in a matching market in practice, decision makers might simply rerun the initially employed matching algorithm (the popular Gale-Shapely algorithm in our case) to compute the new matching M_2 . In Figure 2a, in the dotted line, we indicate the normalized symmetric difference between M_1 , which is the men-optimal matching in \mathcal{P}_1 , and the men-optimal matching in \mathcal{P}_2 . Overall, for all three types of changes and independent of the applied fraction of changes, the normalized symmetric difference between the two men-optimal matchings is close to the minimum achievable normalized symmetric difference, being on average always only at most 0.05 higher (i.e., 5 edges larger) than for the optimal solution.

Since the Gale-Shapley solution has such a good quality, one might conjecture that all stable matchings in \mathcal{P}_2 are roughly similarly different from M_1 . To check this hypothesis, in Figure 2a, in the dashed line, we display the average normalized symmetric difference of M_1 and the stable matching in \mathcal{P}_2 that is furthest away from M_1 . For **Delete**, the above hypothesis actually gets confirmed: after few changes, the worst, the men-optimal, and the best stable matching in \mathcal{P}_2 have a similar distance to M_1 , indicating that after randomly deleting some agents it does not really matter which stable matching in \mathcal{P}_2 is chosen.⁶ In contrast to this, for the other two types of changes, there is a significant difference between the best and worst solution.

To check whether the above observations also hold if we modify the setup of our experiment, we reran it, but now instead of having random preferences in \mathcal{P}_1 , for each of the two sides all agents have the same preferences in \mathcal{P}_1 . Figure 2b displays the results. Comparing Figure 2a and Figure 2b, there are two major differences. First, for identical preferences few changes have an even stronger effect than for random preferences. Second, for identical preferences in \mathcal{P}_1 , all stable matchings in \mathcal{P}_2 have nearly the same distance to M_1 .

Moreover, using Mallows model (Mallows 1957), we also conducted an analysis of cases in between the two considered extremes (random and identical preferences), i.e., if the preferences of agents have some structure. For **Swap** and **Reorder**, the more unstructured agent’s preferences are, the smaller is the minimum normalized symmetric difference between M_1 and a stable matching in \mathcal{P}_2 . In contrast to this, for **Delete**, this quantity first decreases and then increases again when making preferences more and more random. We also repeated the experiment from this section for different numbers of agents without any major changes in the results.

6.2 Almost Stable Marriage

As featured in Section 4, we now analyze the trade-off between the number of pairs that are allowed to block M_2 and the minimum symmetric difference between M_1 and M_2 .

Experimental Setup. For our three different types of changes, for $r \in \{0, 0.01, 0.02, \dots, 0.3\}$, and for $\beta \in$

⁶On a theoretical level, a possible explanation for this is a result of Ashlagi, Kanoria, and Leshno (2017), who proved that in SM instances with an unequal number of men and women and random preferences, stable matchings are “essentially unique”.

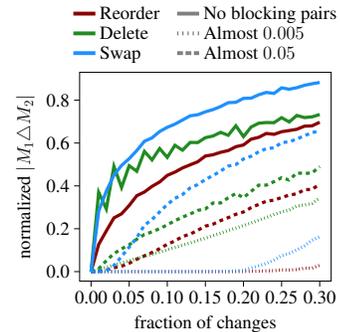


Figure 3: Average normalized symmetric difference between M_1 and a matching M_2 in \mathcal{P}_2 with at most a given number of blocking pairs for a varying fraction of change between \mathcal{P}_1 and \mathcal{P}_2 .

$\{0, 0.005, 0.05\}$, we prepared 200 instances $(U \cup W, \mathcal{P}_1, \mathcal{P}_2)$ as in Section 6.1, where \mathcal{P}_1 and \mathcal{P}_2 differ in an r -fraction of all possible changes. Then, we computed the minimum symmetric difference between M_1 and a matching M_2 in \mathcal{P}_2 for which at most a β -fraction of all $50 \cdot 50$ man-woman pairs is blocking. Figure 3 shows the results of this experiment.⁷

Evaluation. We observe that independent of the type and fraction of change, allowing for few blocking pairs for M_2 enables a significantly larger overlap of M_2 with M_1 . That is, allowing for a 0.005 fraction of pairs to be blocking decreases the average normalized symmetric difference by around 0.2. We also examined the effect of doubling the fraction of blocking pairs and allowing for a 0.01 fraction, which gives an additional decrease by 0.1. If we allow for a 0.05 fraction of pairs to be blocking, then, for **Swap** and **Reorder**, until a 0.2 fraction of changes, M_2 can be chosen to be almost identical to M_1 .

6.3 Further Experiments

We also experimentally explored further aspects of ISM which we briefly summarize here. First, in addition to the three presented types of changes, we considered **Reorder** (inverse) where we reverse the preference list of one agent, and **Add** where we add an agent to the instance. While **Reorder** (inverse) behaves similarly to **Replace** yet typically requiring few more adjustments, **Add** produces results quite similar to **Delete**. Second, we counted blocking pairs for M_1 in \mathcal{P}_2 and observed that their number is highly correlated to the value of an optimal solution, implying that the number of blocking pairs might be used to predict the number of necessary adjustments.

⁷The y -axis here is labeled differently than before: We divide $|M_1 \Delta M_2|$ by the size of a stable matching in \mathcal{P}_1 plus the size of a stable matching in \mathcal{P}_2 . As all stable matchings have the same size, this is the same as $\frac{|M_1 \Delta M_2|}{|M_1| + |M_2|}$ if M_2 is a stable matching; however, different almost stable matchings may have different sizes.

Acknowledgments

NB was supported by the DFG project MaMu (NI 369/19) and by the DFG project ComSoc-MPMS (NI 369/22). KH was supported by the DFG Research Training Group 2434 “Facets of Complexity” and by the DFG project FPTinP (NI 369/16).

References

- Akbarpour, M.; Li, S.; and Gharan, S. O. 2020. Thickness and Information in Dynamic Matching Markets. *J. Political Econ.*, 128(3): 783–815.
- Ashlagi, I.; Kanoria, Y.; and Leshno, J. D. 2017. Unbalanced Random Matching Markets: The Stark Effect of Competition. *J. Political Econ.*, 125(1): 69–98.
- Baccara, M.; Lee, S.; and Yariv, L. 2020. Optimal dynamic matching. *Theor. Econ.*, 15(3): 1221–1278.
- Bhattacharya, S.; Hoefler, 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 (ICALP '15)*, 504–515. Springer.
- Boehmer, N.; Bredereck, R.; Faliszewski, P.; and Niedermeier, R. 2021. Winner Robustness via Swap- and Shift-Bribery: Parameterized Counting Complexity and Experiments. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI '21)*, 52–58. ijcai.org.
- Boehmer, N.; and Heeger, K. 2021. A Fine-Grained View on Stable Many-To-One Matching Problems with Lower and Upper Quotas. *CoRR*, abs/2009.14171v2.
- Boehmer, N.; Heeger, K.; and Niedermeier, R. 2021. Theory of and Experiments on Minimally Invasive Stability Preservation in Changing Two-Sided Matching Markets. *CoRR*, abs/2112.05777.
- Boehmer, N.; and Niedermeier, R. 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)*, 1–5. ACM.
- Bredereck, R.; Chen, J.; Knop, D.; Luo, J.; and Niedermeier, R. 2020. Adapting Stable Matchings to Evolving Preferences. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI '20)*, 1830–1837. AAAI Press. Full version available at arxiv.org/abs/1907.01375.
- Cai, L.; and Thomas, C. 2021. The Short-Side Advantage in Random Matching Markets. *CoRR*, abs/1910.04406v3.
- Cechlárová, K.; and Fleiner, T. 2005. On a generalization of the stable roommates problem. *ACM Trans. Algorithms*, 1(1): 143–156.
- Damiano, E.; and Lam, R. 2005. Stability in dynamic matching markets. *Games Econ. Behav.*, 52(1): 34–53.
- Doval, L. 2021. Dynamically Stable Matching. *CoRR*, abs/1906.11391v5.
- Duan, R.; and Su, H. 2012. A scaling algorithm for maximum weight matching in bipartite graphs. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*, 1413–1424. SIAM.
- Feder, T. 1992. A New Fixed Point Approach for Stable Networks and Stable Marriages. *J. Comput. Syst. Sci.*, 45(2): 233–284.
- Feigenbaum, I.; Kanoria, Y.; Lo, I.; and Sethuraman, J. 2020. Dynamic Matching in School Choice: Efficient Seat Reallocation After Late Cancellations. *Manag. Sci.*, 66(11): 5341–5361.
- Gajulapalli, K.; Liu, J. A.; Mai, T.; and Vazirani, V. V. 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)*, 21:1–21:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Gale, D.; and Shapley, L. S. 1962. College Admissions and the Stability of Marriage. *Amer. Math. Monthly*, 69(1): 9–15.
- Ghosal, P.; Kunysz, A.; and Paluch, K. E. 2020. The Dynamics of Rank-Maximal and Popular Matchings. *CoRR*, abs/1703.10594v3.
- Gupta, S.; Jain, P.; Roy, S.; Saurabh, S.; and Zehavi, M. 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)*, 24:1–24:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Gurobi Optimization, LLC. 2021. Gurobi Optimizer Reference Manual. <https://www.gurobi.com>. Accessed: 2022-04-25.
- Kanade, V.; Leonardos, N.; and Magniez, F. 2016. Stable Matching with Evolving Preferences. In *Proceedings of Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM 2016)*, 36:1–36:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Knuth, D. E.; Motwani, R.; and Pittel, B. G. 1990. Stable Husbands. *Random Struct. Algorithms*, 1(1): 1–14.
- Liu, C. 2021. Stability in Repeated Matching Markets. *CoRR*, abs/2007.03794v2.
- Mallows, C. L. 1957. Non-null ranking models. I. *Biometrika*, 44: 114–130.
- Nimbhorkar, P.; and Rameshwar, V. A. 2019. Dynamic rank-maximal and popular matchings. *J. Comb. Optim.*, 37(2): 523–545.
- Pittel, B. G. 1989. The Average Number of Stable Matchings. *SIAM J. Discret. Math.*, 2(4): 530–549.
- Wilde, H.; Knight, V. A.; and Gillard, J. W. 2020. Matching: A Python library for solving matching games. *J. Open Source Softw.*, 5(48): 2169.