Stable Marriage with Covering Constraints– A Complete Computational Trichotomy

Matthias Mnich[1,2,⋆] and Ildikó Schlotter[3,⋆⋆]

[1] Universität Bonn, Bonn, Germany
[2] Maastricht University, Maastricht, The Netherlands
`mmnich@uni-bonn.de`
[3] Budapest University of Technology and Economics, Budapest, Hungary
`ildi@cs.bme.edu`

Abstract. We consider STABLE MARRIAGE WITH COVERING CONSTRAINTS (SMC): in this variant of STABLE MARRIAGE, we distinguish a subset of women as well as a subset of men, and we seek a matching with fewest number of blocking pairs that matches all of the distinguished people. We investigate how a set of natural parameters, namely the maximum length of preference lists for men and women, the number of distinguished men and women, and the number of blocking pairs allowed determine the computational tractability of this problem.

Our main result is a complete complexity trichotomy that, for each choice of the studied parameters, classifies SMC as polynomial-time solvable, NP-hard and fixed-parameter tractable, or NP-hard and W[1]-hard. We also classify all cases of one-sided constraints where only women may be distinguished.

1 Introduction

The STABLE MARRIAGE (SM) problem is a fundamental problem first studied by Gale and Shapley [16] in 1962. An instance of SM consists of a set $\mathcal{M}$ of men, a set $\mathcal{W}$ of women, and a preference list for each person ordering members of the opposite sex. We aim to find a stable matching, i.e., a matching for which there exists no pair of a man and a woman who prefer each other to their partners given by the matching; such a pair is called a blocking pair.

We consider a problem that we call STABLE MARRIAGE WITH COVERING CONSTRAINTS (SMC). Here, a set $\mathcal{W}^\star$ of women and a set $\mathcal{M}^\star$ of men are distinguished, and a feasible matching is one where each person in $\mathcal{W}^\star \cup \mathcal{M}^\star$ gets matched. By the Rural Hospitals Theorem [17] we know that the set of unmatched men and women is the same in all stable matchings, so clearly, feasible stable matchings may not exist. Thus, we define the task in SMC as finding a feasible matching with a minimum number of blocking pairs. Somewhat surprisingly, this natural extension of SM has not been considered before.

Motivation. Our main motivation for studying SMC—apart from its natural definition—is its close relationship with the HOSPITALS/RESIDENTS WITH LOWER QUOTA (HRLQ) problem, modelling a situation where medical residents apply for jobs in hospitals: residents rank hospitals and vice versa, and hospitals declare both *lower* and *upper quotas* which bound the number of residents they can accept; the task is to find an assignment with a minimum number of blocking pairs. By "cloning" hospitals, HRLQ reduces to the case where each hospital has unit upper quota. In fact, this is equivalent to the special case of SMC where only women (or, equivalently, men) are distinguished. We refer to this problem with one-sided covering constraints, linking SMC and HRLQ, as SMC-1.

The HRLQ problem and its variants have recently gained quite some interest from the algorithmic community [4, 7, 13, 18, 20, 21, 25, 33, 37]. In his book, Manlove [29, Chapter 5.2] devotes an entire chapter to the algorithmics of HRLQ.

The reason for this interest in HRLQ is explained by its importance in several real-world matching markets [14, 15, 35] such as school admission systems, centralized assignment of residents to hospitals, or of cadets to military branches. Lower quotas are a common feature of such admission systems. Their purpose is often to remedy the effects of understaffing that are explained by the well-known Rural Hospitals Theorem [17]: as an example, governments usually want to assign at least a small number of medical residents to each rural hospital to guarantee a minimum service level. Minimum quotas also appear in controlled school choice programs [11, 28, 36] where students belong to a small number of types; schools set lower bounds for each type enact affirmative actions, such as admitting a certain number of minority students [11]. Another example is the German university admission system for admitting students to highly oversubscribed subjects, where a certain percentage of study places are assigned according to high school grades or waiting time [36]. But lower quotas may also arise due to financial considerations: for instance, a business course with too few (tuition-paying) attendees may not be profitable. Certain aspects of airline preferences for seat upgrade allocations can be also modelled by lower quotas [28].

Another motivation for studying SMC comes from the following scenario that we dub *Control for Stable Marriage*. Consider a two-sided market where each participant of the market expresses its preferences over members of the other party, and some central agent (e.g., a government) performs the task of finding a stable matching in the market. It might happen that this central agency wishes to apply a certain *control* on the stable matching produced: it may favour some participants by trying to assign them a partner in the resulting matching. Such a behaviour might be either malicious (e.g., the central agency may accept bribes and thus favour certain participants) or beneficial (e.g., it may favour those who are at disadvantage, like handicapped or minority participants). However, there might not be a stable matching that covers all participants the agency wants to favour; thus arises the need to produce a matching that is *as stable as possible* among those that fulfill our constraints—the most natural aim in such a case is to minimize the number of blocking pairs in the produced matching, which yields exactly the SMC problem. Similar control problems for voting systems have been

extensively studied in social choice following the work initiated by Bartholdi III. [3], but have not yet been considered in connection to stable matchings.

Our Results. We provide an extensive algorithmic analysis of the SMC problem, and its special case SMC-1. In our analysis, we examine the influence of different aspects on the tractability of these problems. The aspects we consider are

- the number b of blocking pairs allowed,
- the number $|\mathcal{W}^\star|$ of women with covering constraint,
- the number $|\mathcal{M}^\star|$ of men with covering constraint,
- the maximum length $\Delta_{\mathcal{W}}$ of women's preference lists, and
- the maximum length $\Delta_{\mathcal{M}}$ of men's preference lists.

To investigate how these aspects affect the complexity of SMC, we use the framework of parameterized complexity, which deals with computationally hard problems and focuses on how certain *parameters* of a problem instance influence its tractability. We aim to design *fixed-parameter algorithms*, which perform well in practice if the value of the parameter on hand is small[4].

The choice of the above aspects (or parameters) is motivated by the aforementioned applications. For instance, we seek matchings where ideally no blocking pairs at all or at least only few of them appear, to ensure stability of the matching and happiness of those getting matched. The number of women/men with covering constraints corresponds, for instance, to the number of rural hospitals for which a minimum quota specifically must be enforced, which we can expect to be small among the set of all hospitals accepting medical residents. Finally, preference lists of hospitals and residents can be expected to be small, as each hospital might not rank many more candidates than positions it has to fill, whereas residents might rank only their top choices of hospitals. Hence, it is reasonable to assume that these parameters take small values in certain applications, and thus suitable fixed-parameter algorithms may be highly efficient in practice.

We draw a detailed landscape of the influence of each aspect, and *all* their combinations, on the complexity of the SMC problem. To this end, we consider all choices of aspects in $A = \{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_{\mathcal{M}}, \Delta_{\mathcal{W}}\}$ as either being restricted to some constant integer, or regarded as a parameter, or left unbounded. Intuitively, these different choices for elements of A correspond to their expected "range" in applications, from very small to mid-range to large (compared to the size of the entire system). By considering all combinations, we can model all applications.

Our main result classifies the SMC problem for all such combinations into one of three cases, as being either "easy", "moderate", or "hard" to solve:

Theorem 1. *For each choice of aspects in* $A = \{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_{\mathcal{M}}, \Delta_{\mathcal{W}}\}$, *SMC is in* P, *or* NP-*hard and fixed-parameter tractable, or* NP-*hard and* W[1]-*hard with the given parameterization[5], and is covered by one of the results in Table 1.*

In particular, SMC is W[1]-*hard parameterized by* $b + |\mathcal{W}^\star|$, *even if there are no distinguished men (i.e.,* $|\mathcal{M}^\star| = 0$), $\Delta_{\mathcal{M}} = 3$, $\Delta_{\mathcal{W}} = 3$ *and each distinguished woman finds only a single man acceptable.*

---

[4] For background on parameterized complexity, we refer to the recent monograph [9].
[5] Restrictions without parameters are classified as polynomial-time solvable or NP-hard.

Table 1 summarizes our results on the complexity of SMC. Note that some results are implied directly by the symmetrical roles of men and women in SMC, and thus are not stated explicitly. Proofs marked by ($\star$), as well as a decision diagram showing that the presented results indeed cover all restrictions of SMC with respect to $\{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_\mathcal{M}, \Delta_\mathcal{W}\}$, can be found in the full version [32].

constants	parameters	complexity
$|\mathcal{M}^\star| = 0, |\mathcal{W}^\star| = 0$		in P (Gale-Shapley alg.)
$|\mathcal{M}^\star| = 0, |\mathcal{W}^\star|, \Delta_\mathcal{M}$		in P (Thm. 8)
$|\mathcal{M}^\star|, |\mathcal{W}^\star|, \Delta_\mathcal{M}, \Delta_\mathcal{W}$		in P (Thm. 8)
$|\mathcal{M}^\star| = 0, \Delta_\mathcal{M} \leq 2$		in P (Thm. 9)
$\Delta_\mathcal{W} \leq 2, \Delta_\mathcal{M} \leq 2$		in P (Observ. 11)
b		in P (Observ. 5)
$|\mathcal{M}^\star| = 0, \Delta_\mathcal{W} \leq 2, \Delta_\mathcal{M} \geq 3$		NP-hard (Thm. 10)
$|\mathcal{W}^\star| = 1, \Delta_\mathcal{W} \leq 2, \Delta_\mathcal{M} \geq 3$		NP-hard (Thm. 12)
$|\mathcal{M}^\star| = 0, \Delta_\mathcal{W} \geq 3, \Delta_\mathcal{M} \geq 3, \Delta^\star = 1$	$b + |\mathcal{W}^\star|$	W[1]-hard (Thm. 2)
$|\mathcal{M}^\star| = 0, |\mathcal{W}^\star| \geq 1, \Delta_\mathcal{W} \geq 3, \Delta^\star = 1$	$b + \Delta_\mathcal{M}$	W[1]-hard (Thm. 6)
$\Delta_\mathcal{W} \leq 2$	$|\mathcal{W}^\star| + |\mathcal{M}^\star|$	FPT (Thm. 13)
$\Delta_\mathcal{W} \leq 2$	b	FPT (Cor. 2)

Table 1: Summary of our results for STABLE MARRIAGE WITH COVERING CONSTRAINTS. Here, $\Delta^\star$ denotes the maximum length of the preference list of any distinguished person.

As a special case, we answer a question by Hamada et al. [20] who gave an exponential-time algorithm that in time $O(|I|^{b+1})$ decides for a given instance I of HRLQ whether it admits a feasible matching with at most b blocking pairs[6]; the authors asked whether HRLQ is fixed-parameter tractable parameterized by b. As shown by Theorem 1, SMC-1 and therefore also HRLQ is W[1]-hard when parameterized by b, already in a very restricted setting. Thus, the answer to the question by Hamada et al. [20] is negative: SMC-1, and hence HRLQ, admits no fixed-parameter algorithm with parameter b unless FPT = W[1].

Related Work. There is a fast-growing body of literature on matching markets with lower quotas [4, 7, 13–15, 18, 20, 21, 25, 33, 37]. These papers study several variants of HRLQ, adapting the general model to the various specialties of practical problems. However, there are only a few papers which consider the problem of minimizing the number of blocking pairs [14, 20]. The most closely related work to ours is that of Hamada et al. [20]: they prove that the HRLQ problem is NP-hard and give strong inapproximability results; they also consider the SMC-1 problem directly and propose an $O(|I|^{b+1})$ time algorithm for it.

A different line of research connected to SMC is the problem of *arranged marriages*, an early extension of SM suggested by Knuth [26] in 1976. Here, a

---

[6] Hamada et al. claim only a run time $O((|\mathcal{W}||\mathcal{M}|)^{b+1})$, but their algorithm can easily be implemented to run in time $O(|I|^{b+1})$.

set $\mathcal{Q}^\star$ of man-woman pairs is distinguished, and we seek a stable matching that contains $\mathcal{Q}^\star$ as a subset. Thus, as opposed to SMC, we not only require that each distinguished person is assigned *some* partner, but instead prescribe its partner exactly. Initial work on arranged marriages [19, 26] was extended by Dias et al. [10] to consider also *forbidden marriages*, and was further generalized by Fleiner et al. [12] and Cseh and Manlove [8]. Despite the similar flavour, none of these papers have a direct consequence on the complexity of SMC.

Our work also fits the research line that addresses computationally hard stable matching problems by focusing on instances with bounded preference lists [6, 22, 24, 27, 34] or by studying their parameterized complexity [1, 2, 5, 30, 31].

2   Preliminaries

An instance I of the STABLE MARRIAGE (SM) problem consists of a set $\mathcal{M}$ of men and a set $\mathcal{W}$ of women. Each person $x \in \mathcal{M} \cup \mathcal{W}$ has a preference list $L(x)$ that strictly orders the members of the other party acceptable for x. We thus write $L(x)$ as a vector $L(x) = (y_1, \ldots, y_t)$, denoting that y_i is (strictly) *preferred* by x over y_j for each i and j with $1 \leq i < j \leq t$. A *matching* M for I is a set of man-woman pairs appearing in each other's preference lists such that each person is contained in at most one pair of M; some persons may be left unmatched by M. For each person x we denote by $M(x)$ the person assigned by M to x. For a matching M, a man m and a woman w included in each other's preference lists form a *blocking pair* if (i) m is either unmatched or prefers w to $M(m)$, and (ii) w is either unmatched or prefers m to $M(w)$. In the STABLE MARRIAGE WITH COVERING CONSTRAINTS (SMC) problem, we are given additional subsets $\mathcal{W}^\star \subseteq \mathcal{W}$ and $\mathcal{M}^\star \subseteq \mathcal{M}$ of *distinguished* people that must be matched; a matching M is *feasible* if it matches everybody in $\mathcal{W}^\star \cup \mathcal{M}^\star$. The objective of SMC is to find a feasible matching for I with minimum number of blocking pairs. If only people from one gender are distinguished, then w.l.o.g., we assume these to be women; this special case will be denoted by SMC-1.

The many-to-one extension of SMC-1 is the HOSPITALS/RESIDENTS WITH LOWER QUOTAS (HRLQ) problem whose input consists of a set $\mathcal{R}$ of residents and a set $\mathcal{H}$ of hospitals that have ordered preferences over the acceptable members of the other party. Each hospital $h \in \mathcal{H}$ has a quota lower bound $\underline{\mathsf{q}}(h)$ and a quota upper bound $\bar{\mathsf{q}}(h)$. One seeks an *assignment* M that maps a subset of the residents to hospitals that respects acceptability and is *feasible*, that is, $\underline{\mathsf{q}}(h) \leq |M(h)| \leq \bar{\mathsf{q}}(h)$ for each hospital h. Here, $M(h)$ is the set of residents assigned to some $h \in \mathcal{H}$ by M. We say that a hospital h is *under-subscribed* if $|M(h)| < \bar{\mathsf{q}}(h)$. For an assignment M of an instance of HRLQ, a pair $\{r, h\}$ of a resident r and a hospital h is *blocking* if (i) r is unassigned or prefers h to the hospital assigned to r by M, and (ii) h is under-subscribed or prefers r to one of the residents in $M(h)$. The task in HRLQ is to find a feasible assignment with minimum number of blocking pairs.

Some instances of SMC may admit a *master list* over women, which is a total ordering $L_\mathcal{W}$ of all women, such that for each man $m \in \mathcal{M}$, the preference

list $L(m)$ is the restriction of $L_\mathcal{W}$ to those women that m finds acceptable. Similarly, we consider master lists over men.

With each instance I of SMC (or HRLQ) we can naturally associate a bipartite graph G_I whose vertex partitions correspond to $\mathcal{M}$ and $\mathcal{W}$ (or $\mathcal{R}$ and $\mathcal{H}$, respectively), and there is an edge between a man $m \in \mathcal{M}$ and a woman $w \in \mathcal{W}$ (or between a resident $r \in \mathcal{R}$ and a hospital $h \in \mathcal{H}$, respectively) if they appear in each other's preference lists. We may refer to entities of I as vertices, or a pair of entities as edges, without mentioning G_I explicitly.

3 Strong Parameterized Intractability of SMC

This section shows parameterized intractability and inapproximability results for finding feasible matchings with minimum number of blocking pairs. A fundamental hypothesis about the complexity of NP-hard problems is the *Exponential Time Hypothesis* (ETH), which stipulates that algorithms solving all SATISFIABILITY instances in subexponential time cannot exist [23].

Theorem 2 ($\star$). *SMC-1 is* W[1]*-hard parameterized by $b + |\mathcal{W}^\star|$, and cannot be solved in time $f'(b) \cdot n^{o(\sqrt{b})}$ for any computable function f' unless ETH fails, even if there is a master list over men as well as one over women, all preference lists are of length at most 3, and $|L(w)| = 1$ for each woman $w \in \mathcal{W}^\star$.*

4 Polynomial-Time Approximation

Theorem 3 ($\star$). *Let I be an instance I of HRLQ, $\underline{\mathsf{q}}_\Sigma$ the sum of lower quota bounds taken over all hospitals in I, and $\Delta_\mathcal{R}$ the maximum length of residents' preference lists. There is an algorithm that in polynomial time either outputs a feasible assignment for I with at most $(\Delta_\mathcal{R} - 1)\underline{\mathsf{q}}_\Sigma$ blocking pairs, involving only $\underline{\mathsf{q}}_\Sigma$ residents, or concludes that no feasible assignment exists.*

If both $\Delta_\mathcal{R}$ and $\underline{\mathsf{q}}_\Sigma$ are constant, then Thm. 3 implies that HRLQ becomes polynomial-time solvable: if $b \geq (\Delta_\mathcal{R} - 1)\underline{\mathsf{q}}_\Sigma$, then we apply Thm. 3 directly; if $b < (\Delta_\mathcal{R} - 1)\underline{\mathsf{q}}_\Sigma$, then we use the algorithm by Hamada et al. [20] running in time $O(|I|^{b+1})$ which is polynomial, since b is upper-bounded by a constant.

Corollary 1. *HRLQ with both the maximum length $\Delta_\mathcal{R}$ of residents' preference lists and the total sum $\underline{\mathsf{q}}_\Sigma$ of all lower quotas constant, is polynomial-time solvable.*

Another application of Thm. 3 is an approximation algorithm that works regardless of whether $\Delta_\mathcal{R}$ or $\underline{\mathsf{q}}_\Sigma$ is a constant. In fact, the algorithm of Thm. 3 can be turned into a $(\Delta_\mathcal{R} - 1)\underline{\mathsf{q}}_\Sigma$-factor approximation algorithm as follows. First, find a stable assignment M_s for I in polynomial time using the extension of the Gale-Shapley algorithm for the HOSPITALS/RESIDENTS problem. If M_s is not feasible, then by the Rural Hospitals Theorem [17], any feasible assignment for I must admit at least one blocking pair; hence, the algorithm of Thm. 3 yields an approximation with (multiplicative and also additive) factor $(\Delta_\mathcal{R} - 1)\underline{\mathsf{q}}_\Sigma$.

We also state an analogue of Thm. 3 that deals with SMC: it handles covering constraints on both sides, but assumes that all quota upper bounds are 1:

Theorem 4 ($\star$). *There is an algorithm that in polynomial time either outputs a feasible matching for an instance I of SMC with at most $(\Delta_{\mathcal{W}}-1)|\mathcal{M}^\star|+(\Delta_{\mathcal{M}}-1)|\mathcal{W}^\star|$ blocking pairs, or concludes that no feasible matching exists for I.*

5   SMC with Bounded Number of Distinguished Persons or Blocking Pairs

In Thm. 2 we proved $\mathsf{W}[1]$-hardness of SMC-1 for the case where $\Delta_{\mathcal{M}} = \Delta_{\mathcal{W}} = 3$, with parameter $b + |\mathcal{W}^\star|$. Here we investigate those instances of SMC and SMC-1 where the length of preference lists may be unbounded, but either b, or the number of distinguished persons is constant.

First, if the number b of blocking pairs allowed is constant, then SMC can be solved by simply running the extended Gale-Shapley algorithm after guessing and deleting all blocking pairs. This complements the result by Hamada et al. [20].

Observation 5. *SMC can be solved in time $O(|I|^{b+1})$, where b denotes the number of blocking pairs allowed in the input instance I.*

In Thm. 6 we prove hardness of SMC-1 even if only one woman must be covered. If we require preferences to follow master lists, then a slightly weaker version of Thm. 6, where $|\mathcal{W}^\star| = 2$, still holds.

Theorem 6 ($\star$). *SMC-1 is $\mathsf{W}[1]$-hard parameterized by $b + \Delta_{\mathcal{M}}$, even if $\mathcal{W}^\star = \{s\}$, $\Delta_{\mathcal{W}} = 3$, and $|L(s)| = 1$.*

Theorem 7 ($\star$). *SMC-1 is $\mathsf{W}[1]$-hard parameterized by $b + \Delta_{\mathcal{M}}$, even if there is a master list over men as well as one over women, $|\mathcal{W}^\star| = 2$, $\Delta_{\mathcal{W}} \leq 3$, and $|L(w)| = 1$ for each $w \in \mathcal{W}^\star$.*

To contrast our intractability results, we show next that if each of $|\mathcal{W}^\star|$, $|\mathcal{M}^\star|$, $\Delta_{\mathcal{W}}$, and $\Delta_{\mathcal{M}}$ is constant, then SMC becomes polynomial-time solvable. Our algorithm relies on the observation that in this case, the number of blocking pairs in an optimal solution is at most $(\Delta_{\mathcal{M}} - 1)|\mathcal{W}^\star| + (\Delta_{\mathcal{W}} - 1)|\mathcal{M}^\star|$ by Thm. 4. Note that for instances of SMC-1, Thm. 8 yields a polynomial-time algorithm already if both $|\mathcal{W}^\star|$ and $\Delta_{\mathcal{M}}$ are constant.

Theorem 8 ($\star$). *SMC can be solved in time $O(|I|^{(\Delta_{\mathcal{M}}-1)|\mathcal{W}^\star|+(\Delta_{\mathcal{W}}-1)|\mathcal{M}^\star|+1})$.*

Importantly, restricting only three of the values $|\mathcal{W}^\star|$, $|\mathcal{M}^\star|$, $\Delta_{\mathcal{W}}$, and $\Delta_{\mathcal{M}}$ to be constant does not yield tractability for SMC, showing that Thm. 8 is tight. Indeed, Thm. 6 implies that restricting the maximum length of the preference lists on only one side still results in a hard problem: SMC remains $\mathsf{W}[1]$-hard with parameter $b + \Delta_{\mathcal{M}}$, even if $\Delta_{\mathcal{W}} = 3$, $|\mathcal{W}^\star| = 1$, and $|\mathcal{M}^\star| = 0$. Similarly, Thm. 2 shows that the problem remains hard even if $\Delta_{\mathcal{W}} = \Delta_{\mathcal{M}} = 3$ and $|\mathcal{M}^\star| = 0$.

6   SMC with Preference Lists of Length at most Two

We show that the restriction of SMC where the maximum length of preference lists is bounded by 2 on one side leads to polynomial-time algorithms and fixed-parameter algorithms for various parameterizations.

Let I be an instance of SMC with underlying graph G. Let M_s be a stable matching in I, and let $\mathcal{M}_0^\star$ and $\mathcal{W}_0^\star$ denote the set of distinguished men and women, respectively, unmatched by M_s. Furthermore, let $\mathcal{M}_0$ and $\mathcal{W}_0$ denote the set of all men and women, respectively, unmatched by M_s. A path P in G is called an *augmenting path*, if $M_s \Delta P$ is a matching, and either both endpoints of P are in $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$, or one endpoint of P is in $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$, and its other endpoint is not distinguished. We will call an augmenting path P *masculine* or *feminine* if it contains a man in $\mathcal{M}_0^\star$ or a woman in $\mathcal{W}_0^\star$, respectively; if P is both masculine and feminine, then we call it *neutral*. If P is not neutral, we say that it *starts* at the (unique) person from $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$ it contains, and *ends* at its other endpoint.

Covering constraints on one side. We give a polynomial-time algorithm for SMC-1 when each man finds at most two women acceptable, and show NP-hardness of SMC-1 even if each woman finds at most two men acceptable.

Theorem 9. *There is a polynomial-time algorithm for the special case of SMC-1 where each man finds at most two women acceptable.*

The main observation behind Thm. 9 is that if $\Delta_\mathcal{M} \leq 2$, then any two augmenting paths starting from different women in $\mathcal{W}_0^\star$ can only intersect at their endpoints. Thus, we can modify the stable matching M_s by selecting augmenting paths starting from each woman in $\mathcal{W}_0^\star$ in an almost independent fashion: intuitively, we simply need to take care not to choose paths sharing an endpoint—a task which can be managed by finding a bipartite matching in certain auxiliary graph G'. To ensure that the number of blocking pairs in the output is minimized, we will assign costs to the augmenting paths. The cost of an augmenting path P roughly determines the number of blocking pairs introduced when modifying M_s along P (though certain special edges need not be counted); hence, our problem reduces to finding a min-weight bipartite matching in G'.

To present the algorithm of Thm. 9 in detail, we start with the following properties of augmenting paths which are easy to prove assuming that $\Delta_\mathcal{M} \leq 2$:

Proposition 1. *Let P_1 and P_2 be augmenting paths starting at w_1 and w_2, resp. If $w_1 \neq w_2$, then P_1 and P_2 are either vertex-disjoint, or they both end at some $m \in \mathcal{M}_0$, with $V(P_1) \cap V(P_2) = \{m\}$. If there is an edge $\{m, w\}$ of G (with $m \in \mathcal{M}$ and $w \in \mathcal{W}$) connecting P_1 and P_2, then $m \in \mathcal{M}_0$ and P_1 or P_2 must end at m. If $w_1 = w_2$ and P is the maximal common subpath of P_1 and P_2 starting at w_1, then either $V(P_1) \cap V(P_2) = V(P)$, or P_1 and P_2 both end at some $m \in \mathcal{M}_0$ and $V(P_1) \cap V(P_2) = V(P) \cup \{m\}$.*

With a set P of edges (typically a set of augmenting paths) where $M_s \triangle P$ is a matching, we associate a *cost*, which is the number of blocking pairs that $M_s \triangle P$ admits. A pair $\{m, w\}$ for some $m \in \mathcal{M}$ and $w \in \mathcal{W}$ is *special*, if $m \in \mathcal{M}_0$ and w

is the second (less preferred) woman in $L(m)$. As it turns out, such edges can be ignored during certain steps of the algorithm; thus, we let the *special cost* of P be the number of non-special blocking pairs in $M_s \triangle P$.

Lemma 1 ($\star$). *For vertex-disjoint augmenting paths P_1 and P_2 with costs c_1 and c_2, resp., the cost of $P_1 \cup P_2$ is at most $c_1 + c_2$. Further, if the cost of $P_1 \cup P_2$ is less than $c_1 + c_2$, then the following holds for $\{i_1, i_2\} = \{1, 2\}$: there is a special edge $\{m, w\}$ with P_{i_1} ending at m and w appearing on P_{i_2}; moreover, $\{m, w\}$ is blocking in $M_s \triangle P_{i_2}$, but not in $M_s \triangle (P_1 \cup P_2)$.*

We are ready to provide the algorithm, in a sequence of four steps.

Step 1: Computing all augmenting paths. By Prop. 1, if we delete $\mathcal{M}_0$ from the union of all augmenting paths starting at some $w \in \mathcal{W}_0^\star$, then we obtain a tree. Furthermore, these trees are mutually vertex-disjoint for different starting vertices of $\mathcal{W}_0^\star$. This allows us to compute all augmenting paths in linear time, e.g., by an appropriately modified version of the DFS algorithm (so that only augmenting paths are considered). During this process, we can also compute the special cost of each augmenting path in a straightforward way.

Step 2: Constructing an auxiliary graph. Using the results of the computation of Step 1, we construct an edge-weighted single bipartite graph G_{path} as follows. The vertex set of G_{path} is the union of $\mathcal{W}_0^\star$ and $\mathcal{M}_0 \cup \{w' \mid w \in \mathcal{W}_0^\star\}$, so for each woman $w \in \mathcal{W}_0^\star$ we create a corresponding new vertex w'. We add an edge between $w \in \mathcal{W}_0^\star$ and $m \in \mathcal{M}_0$ with weight c if there exists an augmenting path with endpoints w and m having special cost c (and no such path with lower special cost exists). Further, for each $w \in \mathcal{W}_0^\star$ we compute the minimum special cost $c_w^{\min}$ of any augmenting path starting at w and *not* ending in $\mathcal{M}_0$, and add an edge between w and w' with weight $c_w^{\min}$ in G_{path}.

Step 3: Computing a minimum weight matching. We compute a matching M_P in G_{path} covering $\mathcal{W}_0^\star$ and having minimum weight. Observe that such a matching corresponds to a set of augmenting paths $\mathcal{P} = \{P_w \mid w \in \mathcal{W}_0^\star\}$ that are mutually vertex-disjoint by Prop. 1. Recall that the special cost of P_w is the weight of the edge in M_P incident to w.

Step 4: Eliminating blocking special edges. In this step, we modify $\mathcal{P}$ iteratively. We start by setting $\mathcal{P}_{\mathrm{act}} = \mathcal{P}$. At each iteration we modify $\mathcal{P}_{\mathrm{act}}$ as follows. We check whether there exists a special edge $\{m^*, w^*\}$ that is blocking in $M_s \triangle \mathcal{P}_{\mathrm{act}}$. If yes, then notice that m^* is not matched in $M_s \triangle \mathcal{P}_{\mathrm{act}}$, because $\{m^*, w^*\}$ is special and thus $m^* \in \mathcal{M}_0$. Let P be the path of $\mathcal{P}_{\mathrm{act}}$ containing w^*. We modify $\mathcal{P}_{\mathrm{act}}$ by truncating P to its subpath between its starting vertex and w^*, and appending to it the edge $\{m^*, w^*\}$. This way, $\{m^*, w^*\}$ becomes an edge of the matching $M_s \triangle \mathcal{P}_{\mathrm{act}}$. The iteration stops when there is no special edge blocking $M_s \triangle \mathcal{P}_{\mathrm{act}}$. Note that once a special edge ceases to be blocking in $M_s \triangle \mathcal{P}_{\mathrm{act}}$, it cannot become blocking again during this process, so the algorithm performs at most $|\mathcal{M}_0|$ iterations. For each $w \in \mathcal{W}_0^\star$, let P_w^* denote the augmenting path in $\mathcal{P}_{\mathrm{act}}$ covering w at the end of Step 4; we define $\mathcal{P}^* = \{P_w^* \mid w \in \mathcal{W}_0^\star\}$ and output the matching $M_s \triangle \mathcal{P}^*$.

This completes the description of the algorithm; we now provide its analysis.

Lemma 2 ($\star$). $M_{\mathrm{sol}} := M_s \triangle \mathcal{P}^*$ *is a feasible matching for I, and the number of blocking pairs for M_{sol} is at most the weight of M_P.*

To show that our algorithm is correct and M_{sol} is optimal, by Lemma 2 it suffices to prove that the weight of M_P is at most the number of blocking pairs in M^{opt}, an optimal solution in I. To this end, we define a matching covering $\mathcal{W}_0^\star$ in G_{path} whose weight is at most the number of blocking pairs in M^{opt}.

Clearly, $M_s \triangle M^{\mathrm{opt}}$ contains an augmenting path Q_w covering w for each $w \in \mathcal{W}_0^\star$. If some Q_w ends at a man $m \in \mathcal{M}_0$, then clearly no other path in $M_s \triangle M^{\mathrm{opt}}$ can end at m. Take the matching M_Q in G_{path} that includes all pairs $\{m, w\}$ where Q_w ends at $m \in \mathcal{M}_0$ for some $w \in \mathcal{W}_0^\star$. Also, we put $\{w, w'\}$ into M_Q if Q_w does not end at a man of $\mathcal{M}_0$. Note that M_Q is indeed a matching.

It remains to show that the weight of M_Q is at most the number of blocking pairs in M^{opt}. By definition, the weight of M_Q is at most the sum of the special costs of the paths Q_w for every $w \in \mathcal{W}_0^\star$. By Lemma 1, any non-special blocking pair in $M_s \triangle Q_w$ remains a blocking pair in $M_s \triangle (\bigcup_{w \in \mathcal{W}_0^\star} Q_w)$, and hence in M^{opt} as well. Hence, there is a matching in G_{path} with weight at most the number of blocking pairs in an optimal solution, implying the correctness of our algorithm. As the algorithm runs in polynomial time, Thm. 9 follows.

Contrasting Thm. 9, if men may have preference lists of length 3, SMC-1 (and hence SMC) is NP-hard even if each woman finds at most two men acceptable.

Theorem 10 ($\star$). *SMC-1 is NP-hard even if $\Delta_{\mathcal{W}} = 2$ and $\Delta_{\mathcal{M}} = 3$.*

Covering constraints on both sides. If $\max(\Delta_{\mathcal{W}}, \Delta_{\mathcal{M}}) \leq 2$, the graph underlying the instance is a collection of paths and cycles, and therefore:

Observation 11. *SMC with $\max(\Delta_{\mathcal{W}}, \Delta_{\mathcal{M}}) \leq 2$ is polynomial-time solvable.*

Recall that the case where $\Delta_{\mathcal{W}} = 2$ and $\Delta_{\mathcal{M}} = 3$ is NP-hard by Thm. 10, even if there are no distinguished men to be covered. However, switching the role of men and women, Thm. 9 shows that if there are no women to be covered, then $\Delta_{\mathcal{W}} \leq 2$ guarantees polynomial-time solvability for SMC. This raises the natural question whether SMC with $\Delta_{\mathcal{W}} \leq 2$ can be solved efficiently if the number of distinguished women is bounded. Next we show that this is unlikely, as the problem turns out to be NP-hard for $|\mathcal{W}^\star| = 1$.

Theorem 12 ($\star$). *SMC is NP-hard, even if $\Delta_{\mathcal{W}} = 2$, $\Delta_{\mathcal{M}} = 3$ and $|\mathcal{W}^\star| = 1$.*

Contrasting Thm. 12, we establish fixed-parameter tractability of the case $\Delta_{\mathcal{W}} \leq 2$. The relevant cases (whose tractability or intractability does not follow from our results obtained so far) are as follows (assuming $\Delta_{\mathcal{W}} \leq 2$ throughout). First, we can take the number of distinguished persons as parameter (note that we know NP-hardness of the cases where $|\mathcal{W}^\star| = 1$ or $|\mathcal{M}^\star| = 0$). Second, we can consider the number of blocking pairs as the parameter. We show the following:

Theorem 13 ($\star$). *There is a fixed-parameter algorithm for the special case of SMC where each woman finds at most two men acceptable (i.e., $\Delta_{\mathcal{W}} \leq 2$), with parameter the number $|\mathcal{W}_0^\star| + |\mathcal{M}_0^\star|$ of distinguished men and women left unmatched by some stable matching (and hence by any stable matching).*

As each augmenting path contains at least one edge that blocks M^{opt}, the number of blocking pairs admitted by M^{opt} is at least $(|\mathcal{W}_0^\star| + |\mathcal{M}_0^\star|)/2$. Thus:

Corollary 2 ($\star$). *There is a fixed-parameter algorithm with parameter b for the special case of SMC where each woman finds at most two men acceptable.*

References

[1] A. Arulselvan, Á. Cseh, M. Groß, D. F. Manlove, and J. Matuschke. Matchings with lower quotas: Algorithms and complexity. *Algorithmica*, 2016.

[2] H. Aziz, H. G. Seedig, and J. K. von Wedel. On the susceptibility of the deferred acceptance algorithm. In *Proc. AAMAS 2015*, pages 939–947, 2015.

[3] J. J. Bartholdi, III, C. A. Tovey, and M. A. Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8–9):27–40, 1992.

[4] P. Biró, T. Fleiner, R. W. Irving, and D. F. Manlove. The college admissions problem with lower and common quotas. *Theoret. Comput. Sci.*, 411(34-36):3136–3153, 2010.

[5] P. Biró, R. W. Irving, and I. Schlotter. Stable matching with couples: An empirical study. *ACM J. Exp. Algorithmics*, 16, 2011. Article 1.2.

[6] P. Biró, D. F. Manlove, and E. J. McDermid. "Almost stable" matchings in the roommates problem with bounded preference lists. *Theoret. Comput. Sci.*, 432:10–20, 2012.

[7] K. Cechlárová and T. Fleiner. Pareto optimal matchings with lower quotas. In *Proc. COMSOC 2016*, 2016.

[8] Á. Cseh and D. F. Manlove. Stable marriage and roommates problems with restricted edges: complexity and approximability. *Discrete Optim.*, 20:62–89, 2016.

[9] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized algorithms*. Springer, 2015.

[10] V. M. Dias, G. D. da Fonseca, C. M. de Figueiredo, and J. L. Szwarcfiter. The stable marriage problem with restricted pairs. *Theoret. Comput. Sci.*, 306(1–3):391–405, 2003.

[11] L. Ehlers, I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. School choice with controlled choice constraints: Hard bounds versus soft bounds. *J. Econ. Theory*, 153:648–683, 2014.

[12] T. Fleiner, R. W. Irving, and D. F. Manlove. Efficient algorithms for generalised stable marriage and roommates problems. *Theor. Comput. Sci.*, 381:162–176, 2007.

[13] T. Fleiner and N. Kamiyama. A matroid approach to stable matchings with lower quotas. *Math. Oper. Res.*, 41(2):734–744, 2016.

[14] D. Fragiadakis, A. Iwasaki, P. Troyan, S. Ueda, and M. Yokoo. Strategyproof matching with minimum quotas. *ACM Trans. Econ. Comput.*, 4(1), 2016.

[15] D. Fragiadakis and P. Troyan. Improving matching under hard distributional constraints. *Theoret. Econ.*, 12:864–908, 2017.

[16] D. Gale and L. S. Shapley. College Admissions and the Stability of Marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.

[17] D. Gale and M. Sotomayor. Some remarks on the stable matching problem. *Discrete Appl. Math.*, 11(3):223–232, 1985.

[18] M. Goto, A. Iwasaki, Y. Kawasaki, R. Kurata, Y. Yasuda, and M. Yokoo. Strategyproof matching with regional minimum and maximum quotas. *Artificial Intelligence*, 235:40–57, 2016.

[19] D. Gusfield and R. W. Irving. *The stable marriage problem: structure and algorithms.* MIT Press, Cambridge, MA, 1989.

[20] K. Hamada, K. Iwama, and S. Miyazaki. The hospitals/residents problem with lower quotas. *Algorithmica*, 74(1):440–465, 2016.

[21] C.-C. Huang. Classified stable matching. In *Proc. SODA 2010*, pages 1235–1253, 2010.

[22] N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proc. SODA 2005*, pages 53–62, 2005.

[23] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. System Sci.*, 63(4):512–530, 2001.

[24] R. W. Irving, D. Manlove, and G. O'Malley. Stable marriage with ties and bounded length preference lists. *J. Discrete Algorithms*, 7:213–219, 2009.

[25] N. Kamiyama. A note on the serial dictatorship with project closures. *Oper. Res. Lett.*, 41(5):559–561, 2013.

[26] D. E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires.* Les Presses de l'Université de Montréal, Montreal, Que., 1976.

[27] F. Kojima, P. A. Pathak, and A. E. Roth. Matching with couples: Stability and incentives in large markets. *Quarterly J. Econ.*, 128(4):1585–1632, 2013.

[28] S. D. Kominers and T. Sönmez. Matching with slot-specific priorities: Theory. *Theoret. Econ.*, 11(2):683–710, 2016.

[29] D. F. Manlove. *Algorithmics of matching under preferences*, volume 2 of *Series on Theoretical Computer Science.* World Scientific, Singapore, 2013.

[30] D. Marx and I. Schlotter. Parameterized complexity and local search approaches for the stable marriage problem with ties. *Algorithmica*, 58(1):170–187, 2010.

[31] D. Marx and I. Schlotter. Stable assignment with couples: Parameterized complexity and local search. *Discrete Optim.*, 8:25–40, 2011.

[32] M. Mnich and I. Schlotter. Stable marriage with covering constraints: A complete computational trichotomy, 2017. `https://arxiv.org/abs/1602.08230`.

[33] D. Monte and N. Tumennasan. Matching with quorums. *Econ. Lett.*, 120(1):14–17, 2013.

[34] A. E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *Amer. Econ. Rev.*, 89:748–780, 1999.

[35] T. Veskioja. *Stable Marriage Problem and College Admission.* PhD thesis, Department of Informatics, Tallinn University of Technology, 2005.

[36] A. Westkamp. An analysis of the German university admissions system. *Econ. Theory*, 53(3):561–589, 2013.

[37] Y. Yokoi. A generalized polymatroid approach to stable matchings with lower quotas. *Math. Oper. Res.*, 42(1):238–255, 2017.