Finding Stable Matchings That Are Robust to Errors in the Input

Tung Mai
Georgia Institute of Technology, Atlanta, GA, USA
tung.mai@cc.gatech.edu

Vijay V. Vazirani
University of California, Irvine, Irvine, CA, USA
vazirani@ics.uci.edu

Abstract

In this paper, we introduce the issue of finding solutions to the stable matching problem that are robust to errors in the input and we obtain the first algorithmic results on this topic. In the process, we also initiate work on a new structural question concerning the stable matching problem, namely finding relationships between the lattices of solutions of two "nearby" instances.

Our main algorithmic result is the following: We identify a polynomially large class of errors, D, that can be introduced in a stable matching instance. Given an instance A of stable matching, let B be the instance that results after introducing one error from D, chosen via a discrete probability distribution. The problem is to find a stable matching for A that maximizes the probability of being stable for B as well. Via new structural properties of the type described in the question stated above, we give a polynomial time algorithm for this problem.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Stable Matching, Robust

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.60

Funding Supported in part by NSF grant CCF-1815901.

Acknowledgements We wish to thank David Eppstein, Mike Goodrich and Sandy Irani for interesting discussions that sparked off the question addressed in this paper.

1 Introduction

Ever since its introduction in the seminal 1962 paper of Gale and Shapley [3], the stable matching problem has been the subject of intense study from numerous different angles in many fields, including computer science, mathematics, operations research, economics and game theory, e.g., see the books [9, 6, 11]. The very first matching-based market, namely matching medical interns to hospitals, was built around this problem, e.g., see [6, 12]. Eventually, this led to an entire inter-disciplinary field, namely matching and market design [12]. The stable matching problem and market design were the subject of the 2012 Nobel Prize in Economics, awarded to Roth and Shapley [13].

The current paper initiates the study of this problem from yet another angle, namely robustness to errors in the input. To the best of our knowledge, this issue has not been studied in the context of this problem (see also Section 1.2) even though the design of algorithms that produce robust solutions is already a very well established field, especially as pertaining to robust optimization, e.g., see the books [2, 1].

26th Annual European Symposium on Algorithms (ESA 2018).
Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 60; pp. 60:1–60:11
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A particularly impressive aspect of the stable matching problem is its deep and pristine combinatorial structure. This in turn has led to efficient algorithms for numerous questions studied about this problem, e.g., see the books mentioned above. A second major contribution of our paper is initiating work on a new structural question, namely finding relationships between the lattices of solutions of two "nearby" instances. In the current paper, and our followup work [10], we restrict ourselves to "nearby" instances which differ in only one agent's preference list. Clearly, this is only the tip of the iceberg as far as "nearby" instances go. Moreover, the structural results are so clean and extensive that they are likely to find algorithmic applications beyond the problem of finding robust solutions. In particular, with ever more interesting matching-based markets being designed and launched on the Internet [12], these new structural properties could find interesting applications and are worth studying further.

We will introduce the problem of finding robust stable matchings via the following model: Alice has an instance A of the stable matching problem, over n boys and n girls, which she sends it to Bob over a channel that can introduce errors. Let B denote the instance received by Bob. Let D denote a polynomial sized domain from which errors are introduced by the channel; we will assume that the channel introduces at most one error from D. We are also given the discrete probability distribution, p over D, from which the channel picks one error. In addition, Alice sends to Bob a matching, M, of her choice, that is stable for instance A. Since M consists of only $O(n)$ numbers of $O(\log n)$ bits each, as opposed to A which requires $O(n^2)$ numbers, Alice is able to send it over an error-free channel. Now Alice wants to pick M in such a way that it has the highest probability of being stable in the instance received by Bob. Hence she picks M from the set

$$\arg\max_N \{Pr_p[N \text{ is stable for instance } B \mid N \text{ is stable for instance } A]\},$$

We will say that such a matching M is *robust*. We seek a polynomial time algorithm for finding such a matching.

Clearly, the domain of errors, D, will have to be well chosen to solve this problem. A natural set of errors is *simple swaps*, under which the positions of two adjacent boys in a girl's list, or two adjacent girls in a boy's list, are interchanged. We will consider a generalization of this class of errors, which we call *shift*. For a girl g, assume her preference list in instance A is $\{\ldots, b_1, b_2, \ldots, b_k, b, \ldots\}$. Move up the position of b so g's list becomes $\{\ldots, b, b_1, b_2, \ldots, b_k, \ldots\}$, and let B denote the resulting instance. Then we will say that B is obtained from A by a shift. An analogous operation is defined on a boy b's list. The domain D consists of all such shifts; clearly, D is polynomially bounded. We prove the following theorem.

▶ **Theorem 1.** *There is a polynomial time algorithm which given an instance A of the stable matching problem and a probability distribution p over the domain, D, of errors defined above, finds a robust stable matching in A.*

1.1 Overview of results and technical ideas

We first summarize some well-known structural facts, e.g., see [6]. The set of stable matchings of an instance form a distributive lattice: given two stable matchings M and M', their meet and join involve taking, for each boy, the optimal or pessimal choice, respectively. It is easy to show that the resulting two matchings are also stable. The extreme matchings of this lattice are called *boy optimal* and *girl-optimal* matchings. A deep notion about this lattice is that of a rotation. A *rotation*, on an ordered list of k boy-girl pairs, when applied to a matching M in which all these boy-girl pairs are matched to each other, matches each boy to the next girl on the list, closing the list under rotation. The k pairs and the order among them are so chosen that the resulting matching is also stable; moreover, a rotation on a subset of these k pairs, under any ordering, leads to a matching that is not stable. Hence, a rotation can be viewed as a minimal change to the current matching that results in a stable matching. Rotations help traverse the lattice from the boy-optimal to the girl-optimal matching along all possible paths available.

For the given instance, a partial order Π is defined on a subset of rotations; the closed sets of Π are in one-to-one correspondence with the set of stable matchings of the instance. Moreover, if S is such a closed set, then starting in the lattice from the boy-optimal matching and applying the rotations in set S, we reach the stable matching corresponding to S.

Let A and B be two instances of stable matching over n boys and n girls, with sets of stable matchings $\mathcal{M}_A$ and $\mathcal{M}_B$, respectively, and lattices $\mathcal{L}_A$ and $\mathcal{L}_B$, respectively. Then, it is easy to see that the matchings in $\mathcal{M}_A \cap \mathcal{M}_B$ form a sublattice in each of the two lattices. Next assume that instance B results from applying a shift operation, defined above, to instance A. Then, we show that $\mathcal{M}_{AB} = \mathcal{M}_A \setminus \mathcal{M}_B$ is also a sublattice of $\mathcal{L}_A$. We use this fact crucially to show that there is at most one rotation, ρ_{in}, that leads from $\mathcal{M}_A \cap \mathcal{M}_B$ to $\mathcal{M}_{AB}$ and at most one rotation, ρ_{out} that leads from $\mathcal{M}_{AB}$ to $\mathcal{M}_A \cap \mathcal{M}_B$. Moreover, we can obtain efficiently this pair of rotations for each of the polynomially many instances that result from the polynomially many shifts.

It is easy to see that a matching M corresponding to a closed set S is stable in instance B iff whenever $\rho_{\text{in}} \in S$, $\rho_{\text{out}} \in S$. We next give an integer program whose optimal solution is a robust stable matching for the given probability distribution on shifts. The IP has one indicator variable, y_ρ, corresponding to each rotation ρ in Π. The constraints of the program ensure that the set S of rotations that are set to 0 form a closed set. The rest of the constraints and the objective function ensure that the corresponding matching maximizes the probability that it is stable in the erroneous instance B. Finally, we show that the LP-relaxation of this IP always has integral solutions. Hence we obtain a polynomial time algorithm for finding a robust stable matching.

1.2 A matter of nomenclature

Assigning correct nomenclature to a new issue under investigation is clearly critical for ease of comprehension. In this context we wish to mention that very recently, Genc et. al. [4] defined the notion of an (a, b)-supermatch as follows: this is a stable matching in which if any a pairs break up, then it is possible to match them all off by changing the partners of at most b other pairs, so the resulting matching is also stable. They showed that it is NP-hard to decide if there is an (a, b)-supermatch. They also gave a polynomial time algorithm for a very restricted version of this problem, namely given a stable matching and a number b, decide if it is a $(1, b)$-supermatch. Observe that since the given instance may have exponentially many stable matchings, this does not yield a polynomial time algorithm even for deciding if there is a stable matching which is a $(1, b)$-supermatch for a given b.

Genc. et. al. [5] also went on to defining the notion of the most robust stable matching, namely a $(1, b)$-supermatch where b is minimum. We would like to point out that "robust" is a misnomer in this situation and that the name "fault-tolerant" is more appropriate. In the literature, the latter is used to describe a system which continues to operate even in the event of failures and the former is used to describe a system which is able to cope with erroneous inputs, e.g., see the following pages from Wikipedia [15, 14].

2 Preliminaries

2.1 The stable matching problem

The stable matching problem takes as input a set of boys $B = \{b_1, b_2, \dots, b_n\}$ and a set of girls $G = \{g_1, g_2, \dots, g_n\}$; each person has a complete preference ranking over the set of opposite sex. The notation $b_i <_g b_j$ indicates that girl g strictly prefers b_j to b_i in her preference list. Similarly, $g_i <_b g_j$ indicates that the boy b strictly prefers g_j to g_i in his list.

A matching M is a one-to-one correspondence between B and G. For each pair $bg \in M$, b is called the partner of g in M (or M-partner) and vice versa. For a matching M, we say that b is *above* (or *below*) g if he prefers his M-partner to g (or g to his M-partner). Similarly, g is said to be *above* (or *below*) b if she prefers her M-partner to b (or b to her M-partner). For a matching M, a pair $bg \notin M$ is said to be *blocking* if b is below g and g is below b, i.e., they prefer each other to their partners. A matching M is *stable* if there is no blocking pair in M.

2.2 The lattice of stable matchings

Let M and M' be two stable matchings. We say that M *dominates* M', denoted by $M \preceq M'$, if every boy weakly prefers his partner in M to M'. It is well known that the dominance partial order over the set of stable matchings forms a distributive lattice [6], with meet and join defined as follows. The *meet* of M and M', $M \wedge M'$, is defined to be the matching that results when each boy chooses his more preferred partner from M and M'; it is easy to show that this matching is also stable. The *join* of M and M', $M \vee M'$, is defined to be the matching that results when each boy chooses his less preferred partner from M and M'; this matching is also stable. These operations distribute, i.e., given three stable matchings M, M', M'',

$$M \vee (M' \wedge M'') = (M \vee M') \wedge (M \vee M'') \;\text{ and }\; M \wedge (M' \vee M'') = (M \wedge M') \vee (M \wedge M'').$$

It is easy to see that the lattice must contain a matching, M_0, that dominates all others and a matching M_z that is dominated by all others. M_0 is called the *boy-optimal matching*, since in it, each boy is matched to his most favorite girl among all stable matchings. This is also the *girl-pessimal matching*. Similarly, M_z is the *boy-pessimal* or *girl-optimal matching*.

2.3 Rotations help traverse the lattice

A crucial ingredient needed to understand the structure of stable matchings is the notion of a rotation, which was defined by Irving [7] and studied in detail in [8]. A rotation takes r matched pairs in a fixed order, say $\{b_0g_0, b_1g_1, \dots, b_{r-1}g_{r-1}\}$ and "cyclically" changes the mates of these $2r$ agents, as defined below, to arrive at another stable matching. Furthermore, it represents a minimal set of pairings with this property, i.e, if a cyclic change is applied on any subset of these r pairs, with any ordering, then the resulting matching has a blocking pair and is not stable. After rotation, the boys' mates weakly worsen and the girls' mates weakly improve. Thus one can traverse from M_0 to M_z by applying a suitable sequence of rotations (specified by the rotation poset defined below). Indeed, this is precisely the purpose of rotations.

Let M be a stable matching. For a boy b let $s_M(b)$ denote the first girl g on b's list such that g strictly prefers b to her M-partner. Let $next_M(b)$ denote the partner in M of girl $s_M(b)$. A *rotation ρ exposed* in M is an ordered list of pairs $\{b_0g_0, b_1g_1, \dots, b_{r-1}g_{r-1}\}$ such that for each i, $0 \le i \le r-1$, b_{i+1} is $next_M(b_i)$, where $i+1$ is taken modulo r. In this paper, we assume that the subscript is taken modulo r whenever we mention a rotation. Notice that a rotation is cyclic and the sequence of pairs can be rotated. M/ρ is defined to be a matching in which each boy not in a pair of ρ stays matched to the same girl and each boy b_i in ρ is matched to $g_{i+1} = s_M(b_i)$. It can be proven that M/ρ is also a stable matching. The transformation from M to M/ρ is called the *elimination* of ρ from M.

Let $\rho = \{b_0g_0, b_1g_1, \ldots, b_{r-1}g_{r-1}\}$ be a rotation. For $0 \le i \le r-1$, we say that ρ *moves* b_i *from* g_i *to* g_{i+1}, and *moves* g_i *from* b_i *to* b_{i-1}. If g is either g_i or is strictly between g_i and g_{i+1} in b_i's list, then we say that ρ *moves* b_i *below* g. Similarly, ρ *moves* g_i *above* b if b is b_i or between b_i and b_{i-1} in g_i's list.

2.4 The rotation poset

A rotation ρ' is said to precede (or dominate) another rotation ρ, denoted by $\rho' \prec \rho$, if ρ' is eliminated in every sequence of eliminations from M_0 to a stable matching in which ρ is exposed. Thus, the set of rotations forms a partial order via this precedence relationship. The partial order on rotations is called *rotation poset* and denoted by Π.

▶ **Lemma 2** ([6], Lemma 3.2.1). *For any boy b and girl g, there is at most one rotation that moves b to g, b below g, or g above b. Moreover, if ρ_1 moves b to g and ρ_2 moves b from g then $\rho_1 \prec \rho_2$.*

A *closed subset* is a subset of the poset such that if an element is in the subset then all of its predecessors are also included. There is a one-to-one relationship between the stable matchings and the closed subsets of Π. Given a closed subset C, the correponding matching M is found by eliminating the rotations starting from M_0 according to the topological ordering of the elements in the subset. We say that C *generates* M.

▶ **Lemma 3** ([6], Lemma 3.3.2). *Π contains $O(n^2)$ rotations and can be computed in polynomial time.*

▶ **Lemma 4** ([6], Theorem 2.5.4). *Every rotation appears exactly once in any sequence of elimination from M_0 to M_z.*

2.5 The notion of shift

In this paper, we will assume that a girl applies a *shift* to one of her preferences as defined below. We will study the structural properties of the resulting instance B.

Let the preference list of girl g in A be $\{\ldots, b_1, b_2, \ldots, b_k, b, \ldots\}$. In B the preference list of g is $\{\ldots, b, b_1, b_2, \ldots, b_k, \ldots\}$. Moreover, all other preference lists are identical in both A and B. We say that B obtained from A by applying a *shift*. We denote $x <^I_z y$ if z prefers y to x in instance I.

3 Structural Results

3.1 The stable matchings in $\mathcal{M}_A \setminus \mathcal{M}_B$ form a sublattice

Let $\mathcal{M}_A$ and $\mathcal{M}_B$ be the sets of all stable matchings under instance A and B respectively. Let $\mathcal{M}_{AB} = \mathcal{M}_A \setminus \mathcal{M}_B$. In other words, $\mathcal{M}_{AB}$ is the set of stable matchings in A that become unstable in B. In this section we show that $\mathcal{M}_{AB}$ forms a lattice. We first prove a simple observation.

▶ **Lemma 5.** *Let $M \in \mathcal{M}_{AB}$. The only blocking pair of M under instance B is bg.*

Proof. Since $M \notin \mathcal{M}_B$, there must be a blocking pair $xy \notin M$ under B. Assume xy is not bg, we will show that xy must also be a blocking pair in A. Let y' be the partner of x and x' be the partner of y in M. Since xy is a blocking pair in B, $x >_y^B x'$ and $y >_x^B y'$. The preference list of x remain unchanged from A to B, so $y >_x^A y'$. Next, we consider two cases:

- If y is not g, the preference list of y does not change. Therefore, $x >_y^A x'$, and hence, xy is also a blocking pair in A.
- If y is g, for all pairs x, x' such that $x >_y^B x'$ and $x \neq b$, we also have $x >_y^A x'$. Therefore, xy is a blocking pair in A.

This contradicts the fact that M is stable under A. ◀

Recall that $b_1 \geq_g b_2 \geq_g \ldots \geq_g b_k$ are k boys right above b in g's list such that the position of b is shifted up to be above b_1 in B. From Lemma 5, we can then characterize the set $\mathcal{M}_{AB}$.

▶ **Lemma 6.** *$\mathcal{M}_{AB}$ is the set of all stable matchings in A that match g to a partner between b_1 and b_k in g's list, and match b to a partner below g in b's list.*

Proof. Assume M is a stable matching in A that contains $b_i g$ for $1 \leq i \leq k$ and bg' such that $g >_b g'$. In B, g prefers b to b_i, and hence bg is a blocking pair. Therefore, M is not stable under B and $M \in \mathcal{M}_{AB}$.

To prove the other direction, let M be a matching in $\mathcal{M}_{AB}$. By Lemma 5, bg is the only blocking pair of M in B. For that to happen, $p_M(b) <_b^B g$ and $p_M(g) <_g^B b$. We will show that $p_M(g) = b_i$ for $1 \leq i \leq k$. Assume not, then $p_M(g) <_g^B b_k$, and hence, $p_M(g) <_g^A b$. Therefore, bg is a blocking pair in A, which is a contradiction. ◀

Let $\mathcal{L}_A$ be the boy-optimal lattice formed by $\mathcal{M}_A$.

▶ **Theorem 7.** *$\mathcal{M}_{AB}$ forms a sublattice of $\mathcal{L}_A$.*

Proof. Assume $\mathcal{M}_{AB}$ is not empty. Let M_1 and M_2 be two matchings in $\mathcal{M}_{AB}$. By Lemma 6, M_1 and M_2 both match g to a partner between b_1 and b_k in g's list, and match b to a partner below g in b's list. Since $M_1 \wedge M_2$ is the matching resulting from having each boy choose the more preferred partner and each girl choose the least preferred partner, $M_1 \wedge M_2$ also belongs to the set characterized by Lemma 6. A similar argument can be applied to the case of $M_1 \vee M_2$. Therefore $\mathcal{M}_{AB}$ form a sublattice of $\mathcal{L}_A$. ◀

3.2 Rotations going into and out of a sublattice

Let M be a stable matching in $\mathcal{M}_A$ and ρ be a rotation exposed in M with respect to instance A. If $M \notin \mathcal{S}$ and $M/\rho \in \mathcal{S}$ for a set $\mathcal{S}$, we say that ρ *goes into* $\mathcal{S}$. Similarly, if $M \in \mathcal{S}$ and $M/\rho \notin \mathcal{S}$, we say that ρ *goes out of* S. Let the set of all rotations going into $\mathcal{S}$ and out of $\mathcal{S}$ be $I_{\mathcal{S}}$ and $O_{\mathcal{S}}$, respectively.

Let $\{b_{i_1}, \ldots b_{i_l}\}$ be the set of possible partners of g in any stable matching in $\mathcal{M}_{AB}$, where $1 \leq i_1 \leq \ldots \leq i_l \leq k$. Let ρ_1 be a rotation moving g to b_{i_l}, ρ_2 be the rotation moving b below g and ρ_3 be a rotation moving g from b_{i_1}. Note that each of ρ_1, ρ_2 and ρ_3 might not exist.

▶ **Lemma 8.** *$I_{\mathcal{M}_{AB}}$ can only contain ρ_1, ρ_2. $O_{\mathcal{M}_{AB}}$ can only contain ρ_3.*

Proof. Consider a rotation $\rho \in I_{\mathcal{M}_{AB}}$. There exists $M \in \mathcal{M}_A \setminus \mathcal{M}_{AB}$ such that $M/\rho \in \mathcal{M}_{AB}$. By Lemma 6, M/ρ matches g to a partner between b_1 and b_k in g's list, and matches b to a partner below g in b's list. Moreover, M either does not contain $b_i g$ for any $1 \leq i \leq k$, or contains bg' where $g' \geq_b g$, or both. If M does not contain $b_i g$ for any $1 \leq i \leq k$, then $\rho = \rho_1$. If M contains bg' where $g' \geq_b g$, then $\rho = \rho_2$.

Consider a rotation $\rho \in O_{\mathcal{M}_{AB}}$. There exists $M \in \mathcal{M}_{AB}$ such that $M/\rho \in \mathcal{M}_A \setminus \mathcal{M}_{AB}$. Again, by Lemma 6, M contains $b_i g$ for $1 \leq i \leq k$ and bg' where $g' <_b g$. Since M dominates M/ρ in the boy optimal lattice, b must prefer g' to his partner in M/ρ. Hence, M/ρ matches b to a partner below g in b's list. Therefore, M/ρ must not contain $b_i g$ for any $1 \leq i \leq k$. It follows that ρ must be ρ_3. ◀

▶ **Lemma 9.** *If both ρ_1 and ρ_2 exist then $\rho_1 \preceq \rho_2$.*

Proof. Assume that $\rho_1 \neq \rho_2$ and there exists a sequence of rotation eliminations, from M_0 to a stable matching M in which ρ_2 is exposed, that does not contain ρ_1. Since ρ_2 moves b below g, g is matched a partner higher than b in her list in M/ρ_2. Therefore, the partner can only be b_{i_l} or a boy higher than b_{i_l} in g's list.

Consider any sequence of rotation eliminations from M/ρ to M_z. In the sequence, the position of g's partner can only go higher in her list. Therefore, ρ_1 cannot be exposed in any matching in the sequence. It follows that ρ_1 is not exposed in a sequence of eliminations from M_0 to M_z, which is a contradiction by Lemma 4. ◀

▶ **Theorem 10.** *There is at most one rotation in $I_{\mathcal{M}_{AB}}$ and at most one rotation in $O_{\mathcal{M}_{AB}}$. Moreover, the rotation in $I_{\mathcal{M}_{AB}}$ must be either ρ_1 or ρ_2, and the rotation in $O_{\mathcal{M}_{AB}}$ must be ρ_3.*

Proof. By Lemma 8, $I_{\mathcal{M}_{AB}}$ can contain at most 2 rotations, namely ρ_1 and ρ_2 if they are distinct. By Lemma 9, if both of them exist, $\rho_1 \preceq \rho_2$. Hence, $I_{\mathcal{M}_{AB}}$ can contain at most one rotation, and it is either ρ_1 or ρ_2.

Again, by Lemma 8, $O_{\mathcal{M}_{AB}}$ can contain at most one rotation, namely ρ_3 if it exists. ◀

By Theorem 10, there is at most one rotation ρ_{in} coming into $\mathcal{M}_{AB}$ and at most one rotation ρ_{out} coming out of $\mathcal{M}_{AB}$.

▶ **Proposition 11.** *ρ_{in} and ρ_{out} can be computed in polynomial time.*

Proof. Since we can compute Π_A efficiently according to Lemma 3, each of ρ_1, ρ_2 and ρ_3 can be computed efficiently.

First we can check possible partners of b and g with respect to instance A. By Lemma 6, $\mathcal{M}_{AB}$ is empty if none of the possible partners of g is between b_1 and b_k in g's list or none of the partners of b is below g in b's list. It follows that both ρ_{in} and ρ_{out} do not exist. Hence we may assume that such a case does not happen.

Suppose ρ_2 exists. If ρ_3 exists and $\rho_3 \preceq \rho_2$, $\mathcal{M}_{AB} = \emptyset$. Otherwise, $\rho_{\text{in}} = \rho_2$, and $\rho_{\text{out}} = \rho_3$ if ρ_3 exists.

Suppose ρ_2 does not exist. If ρ_1 exists, $\rho_{\text{in}} = \rho_1$. If ρ_3 exists, $\rho_{\text{out}} = \rho_3$. ◀

▶ **Lemma 12.** *Let M be a matching in $\mathcal{M}_{AB}$ and S be the corresponding closed subset in Π_A. If ρ_1 exists, S must contain ρ_1. If ρ_2 exists, S must contain ρ_2. If ρ_3 exists, S must not contain ρ_3.*

Proof. If ρ_1 exists, M_0 does not contain $b_i g$ for any $i \in [1,k]$. Since $M \in \mathcal{M}_{AB}$, by Lemma 6 M matches g to a boy between b_1 and b_k in her list. The set of rotations eliminated from M_0 to M must include ρ_1.

If ρ_2 exists, b can not be below g in M_0. Since b is below g in M, by Lemma 6 the set of rotations eliminated from M_0 to M must include ρ_2.

Assume that ρ_3 exists and S contains ρ_3. Since ρ_3 moves g up from b_{i_1}, M can not contain $b_i g$ for any $i \in [1,k]$. This is a contradiction. ◀

3.3 The rotation poset for the sublattice M_{AB}

From the previous section we know that M_{AB} is a sublattice of M_A. In this section we give the rotation poset that generates all stable matchings in the sublattices.

We may assume that $M_{AB} \neq \emptyset$. If ρ_{in} exists, let $\Pi_{\text{in}} = \{\rho \in \Pi_A : \rho \preceq \rho_{\text{in}}\}$ and M_{boy} be the matching generated by Π_{in}. Otherwise, let $M_{\text{boy}} = M_0$. Similarly, let M_{girl} be the matching generated by $\Pi_A \setminus \Pi_{\text{out}}$, where $\Pi_{\text{out}} = \{\rho \in \Pi_A : \rho \succeq \rho_{\text{out}}\}$, if ρ_{out} exists, and $M_{\text{girl}} = M_z$ otherwise.

▶ **Lemma 13.** *M_{boy} is the boy-optimal matching in $\mathcal{M}_{AB}$, and M_{girl} is the girl-optimal matching in $\mathcal{M}_{AB}$.*

Proof. Let M be a matching in $\mathcal{M}_{AB}$ generated by a closed subset $S \subseteq \Pi_A$. By Lemma 12, if ρ_{in} exists, S must contain ρ_{in}. Since Π_{in} is the minimum set containing ρ_{in}, $\Pi_{\text{in}} \subseteq S$. Therefore, $M_{\text{boy}} \preceq M$.

To prove that $M \preceq M_{\text{girl}}$, we show $S \subseteq \Pi_A \setminus \Pi_{\text{out}}$. Assume otherwise, then there exists a rotation $\rho \in S$ such that $\rho \notin \Pi_A \setminus \Pi_{\text{out}}$. It follows that $\rho \in \Pi_{\text{out}}$, and hence $\rho \succeq \rho_{\text{out}}$. Since S contains ρ and S is a closed subset, S must also contain ρ_{out}. This is a contradiction by Lemma 12. ◀

▶ **Theorem 14.** $\Pi_{AB} = \Pi_A \setminus (\Pi_{in} \cup \Pi_{out})$ *is the rotation poset generating $\mathcal{M}_{AB}$.*

Proof. Let M be a matching in $\mathcal{M}_{AB}$ generated by a closed subset $S \subseteq \Pi_A$. Let $S' = S \setminus \Pi_{\text{in}}$. We show that S' is a closed subset of Π_{AB} and eliminating the rotations in S' starting from M_{boy} according to the topological ordering of the elements gives M.

First $S' \cap \Pi_{\text{in}} = \emptyset$ trivially. Since $M \in \mathcal{M}_{AB}$, S does not contain ρ_{out} by Lemma 12. Therefore, S' does not contain ρ_{out}, and $S' \cap \Pi_{\text{out}} = \emptyset$. It follows that S' is a closed subset of Π_{AB}.

Next observe that we can eliminate rotations in S from M_0 by eliminating rotations in Π_{in} first and then eliminating rotations in $S \setminus \Pi_{\text{in}}$. This can be done because Π_{in} is a closed subset of Π_A. Since Π_{in} generates M, the lemma follows. ◀

4 Algorithm for finding a robust stable matching

We now use the structural properties described in Section 3 to give a polynomial time algorithm for finding a robust stable matching. Clearly, the results in Section 3 can be reproduced when we make a shift in a boy's list. Recall from Section 1 that given a discrete probability distribution $\mathcal{D}$ on all possible shifts, a robust stable matching is a stable matching $M \in \mathcal{M}_A$ that minimizes the probability that $M \in \mathcal{M}_{AB}$, where $B \sim \mathcal{D}$.

For a shift B, let ρ^B_{in} and ρ^B_{out} be the rotation going into $\mathcal{M}_{AB}$ and out of $\mathcal{M}_{AB}$ respectively. By Proposition 11, ρ^B_{in} and ρ^B_{out} can be computed efficiently for each B.

By Lemma 3, Π_A can be computed in polynomial time. We create two additional vertices, a source s and a sink t. For a shift B, we may ignore the cases where neither ρ^B_{in} nor ρ^B_{out} exist. In those cases, either $\mathcal{M}_A = \mathcal{M}_B$ or $\mathcal{M}_A \cap \mathcal{M}_B = \emptyset$. Hence, assume that such an instance B does not exist, and $\mathcal{M}_{AB}$ is always a proper non-empty subset of $\mathcal{M}_A$. For a shift B such that ρ^B_{in} does not exist, let $\rho^B_{\text{in}} = s$. Similarly, for a shift B such that ρ^B_{out} does not exist, let $\rho^B_{\text{out}} = t$.

Let p_B be the probability that instance B is chosen according to D. Consider the following integer program:

$$\begin{aligned}
\min \quad & \sum_B x_B p_B \\
\text{s.t.} \quad & y_{\rho_1} \le y_{\rho_2} && \forall \rho_1, \rho_2 : \rho_1 \prec \rho_2 \\
& y_t = 1 \\
& y_s = 0 \\
& x_B \ge y_{\rho^B_{\text{out}}} - y_{\rho^B_{\text{in}}} && \forall B \\
& x_B \ge 0 && \forall B \\
& y_\rho \in \{0, 1\} && \forall \rho \in \Pi_A.
\end{aligned} \tag{IP}$$

▶ **Lemma 15.** *(IP) gives a solution to a robust stable matching.*

Proof. Let $S = \{\rho : y_\rho = 0\}$. The set of constraints:

$$y_{\rho_1} \le y_{\rho_2} \quad \forall \rho_1, \rho_2 : \rho_1 \prec \rho_2$$

guarantees that S is a closed subset.

Notice that $x_B = 1$ if and only if $y_{\rho^B_{\text{out}}} = 1$ and $y_{\rho^B_{\text{in}}} = 0$. This, in turn, happens if and only if the matching generated by S is in $\mathcal{M}_{AB}$.

Therefore, by minimizing $\sum_{e \in E} x_B p_B$, we can find a closed subset that generates a robust stable matching. ◀

▶ **Lemma 16.** *(IP) can be solved in polynomial time.*

Proof. Consider relaxing the constraint $y_\rho \in \{0, 1\}$ to $0 \le y_\rho \le 1$. We show how to round a solution of this natural LP-relaxation of (IP) to have an integral solution of the same objective function. It suffices to just consider $\boldsymbol{y}$ as x_B will always be set to $\max(0, y_{\rho^B_{\text{out}}} - y_{\rho^B_{\text{in}}})$ for any given $\boldsymbol{y}$.

Let $\boldsymbol{y}$ be a fractional optimal solution of the relaxation. Let $1 = a_0 > a_1 > a_2 > \ldots > a_k > a_{k+1} = 0$ be all the possible y-values. Since $\boldsymbol{y}$ is fractional, $k \ge 1$. Denote S_i by the set of all rotations having y-value equal to a_i, where $1 > a_i > 0$.

Let $\mathcal{B}^+$ be the set of instances B such that:

- $x_B = y_{\rho^B_{\text{out}}} - y_{\rho^B_{\text{in}}} > 0$.
- $y_{\rho^B_{\text{out}}} = a_i$.
- $y_{\rho^B_{\text{in}}} \ne a_i$.

Let $\mathcal{B}^-$ be the set of instances B such that:

- $x_B = y_{\rho^B_{\text{out}}} - y_{\rho^B_{\text{in}}} > 0$.
- $y_{\rho^B_{\text{in}}} = a_i$.
- $y_{\rho^B_{\text{out}}} \ne a_i$.

Consider perturbing the y-value of all rotations in S_a by a small amount ϵ:

$$y_\rho \leftarrow y_\rho + \epsilon = a_i + \epsilon \quad \forall \rho \in S_a.$$

Here ϵ is chosen so that $a_i + \epsilon < a_{i-1}$ and $a_i + \epsilon > a_{i+1}$. The net change in the objective function is

$$\sum_{B \in \mathcal{B}^+} \epsilon p_B - \sum_{B \in \mathcal{B}^-} \epsilon p_B = \epsilon \left(\sum_{B \in \mathcal{B}^+} p_B - \sum_{B \in \mathcal{B}^-} p_B \right).$$

We claim that

$$\sum_{B \in \mathcal{B}^+} p_B - \sum_{B \in \mathcal{B}^-} p_B = 0.$$

Assume otherwise, we can pick a sign of ϵ to have a strictly smaller objective function. Since $\sum_{B \in \mathcal{B}^+} p_B - \sum_{B \in \mathcal{B}^-} p_B = 0$, we can choose $\epsilon = a_{i-1} - a_i$ and obtain another optimal solution where the value of k decreases by 1. Keep going until $k = 0$ gives an integral solution. ◀

Finally, Theorem 1 follows from Lemmas 15 and 16.

5 Discussion

As stated in the Introduction, the two main questions on stable matching introduced in this paper are obtaining efficient algorithms for finding solutions that are robust to errors in the input, and the structural question of finding relationships between the lattices of solutions of two "nearby" instances. The current paper and our followup work [10] seem to suggest that both these issues are likely to lead to much work in the future. In particular, the structural results are so clean and extensive that they are likely to find algorithmic applications beyond the problem of finding robust solutions. One possible domain of applications that may be able to exploit these structural properties is matching-based markets, particularly as we are seeing ever more interesting such markets being designed and launched on the Internet, e.g., see [12].

At a more detailed level, the domain D, for which we have obtained our algorithm, is very restrictive and we need to extend it to a larger domain. Our followup paper [10] does this, though it seems more can be done. In particular, are there ways of dealing with two or more errors? Another interesting question is to improve the running time of our algorithm. This looks also quite plausible.

Beyond these questions, pertaining to the most basic of formulations of stable matching, one can study numerous variants and generalizations, such as incomplete preference lists, the stable roommates problem, and matching intern couples to hospitals. Each of these bring their own structural properties and challenges, e.g., see [6, 11].

References

1. A. Ben-Tal, L. El Ghaoui, and A.S. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, October 2009.
2. G. C. Calafiore and L. El Ghaoui. On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130(1), 2006.
3. David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
4. Begum Genc, Mohamed Siala, Barry O'Sullivan, and Gilles Simonin. Finding robust solutions to stable marriage. *arXiv preprint arXiv:1705.09218*, 2017.

5 Begum Genc, Mohamed Siala, Gilles Simonin, and Barry O'Sullivan. On the complexity of robust stable marriage. In *International Conference on Combinatorial Optimization and Applications*, pages 441–448. Springer, 2017.

6 Dan Gusfield and Robert W Irving. *The stable marriage problem: structure and algorithms*. MIT press, 1989.

7 Robert W Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6(4):577–595, 1985.

8 Robert W Irving and Paul Leather. The complexity of counting stable marriages. *SIAM Journal on Computing*, 15(3):655–667, 1986.

9 Donald Ervin Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*. American Mathematical Soc., 1997.

10 Tung Mai and Vijay V. Vazirani. A generalization of Birkhoff's theorem for distributive lattices, with applications to robust stable matchings. In arXiv, 2018.

11 David Manlove. *Algorithmics of Matching Under Preferences*. World Scientific, 2013.

12 Alvin E. Roth. Al Roth's game theory, experimental economics, and market design page, 2016. URL: `http://stanford.edu/~alroth/alroth.html#MarketDesign`.

13 Alvin E. Roth and Lloyd S. Shapley. The 2012 Nobel Prize in Economics, 2012. URL: `https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2012/`.

14 Wikipedia. Fault Tolerence. URL: `https://en.wikipedia.org/wiki/Fault_tolerance`.

15 Wikipedia. Robustness (Computer Science). URL: `https://en.wikipedia.org/wiki/Robustness_(computer_science)`.