

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Exploiting Structure in the Stable Matching Problem

Permalink

<https://escholarship.org/uc/item/6f80v7jc>

Author

Moeller, Daniel Paul

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Exploiting Structure in the Stable Matching Problem

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy

in

Computer Science

by

Daniel Paul Moeller

Committee in charge:

Professor Ramamohan Paturi, Chair
Professor Sanjoy Dasgupta
Professor Massimo Franceschetti
Professor Russell Impagliazzo
Professor Joel Sobel

2016

Copyright

Daniel Paul Moeller, 2016

All rights reserved.

The Dissertation of Daniel Paul Moeller is approved and is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2016

TABLE OF CONTENTS

Signature Page	iii
Table of Contents	iv
List of Figures	vi
List of Tables	vii
Acknowledgements	viii
Vita	ix
Abstract of the Dissertation	x
Chapter 1 Introduction	1
Chapter 2 Background	5
2.1 Basic Stable Matching Concepts	5
2.2 Related Work	6
2.2.1 The Deferred Acceptance Algorithm	6
2.2.2 Complexity of Stable Matching	6
2.2.3 Variations on Stable Matching	7
2.2.4 Strategic Behavior	8
2.2.5 Structure in Stable Matching	9
Chapter 3 Jealousy Graphs	12
3.1 Preliminaries	14
3.1.1 Better Response Algorithm	14
3.1.2 Jealousy Graph and Related Definitions	15
3.2 Structural Results	18
3.2.1 Any Digraph can be a Jealousy Graph	18
3.2.2 Properties of Decompositions	19
3.3 Convergence	21
3.4 Correlated and Intercorrelated Preferences	29
3.5 Conclusion and Open Problems	32
3.6 Acknowledgements	32
Chapter 4 Succinct Preference Models	34
4.1 Preliminaries	38
4.2 Finding Stable Matchings	39
4.2.1 Small Set of Attributes and Weights	39
4.2.2 One-Sided Real Attributes	40

4.2.3	Strategic Behavior	47
4.3	Verification	48
4.3.1	Real Attributes and Weights	48
4.3.2	Lists	50
4.3.3	Boolean Attributes and Weights	52
4.4	Conditional Hardness	55
4.4.1	Background	55
4.4.2	Finding Stable Matchings	57
4.4.3	Verifying Stable Matchings	58
4.4.4	Checking a Stable Pair	60
4.5	Other Succinct Preference Models	66
4.5.1	One Dimensional Single-Peaked Preferences	68
4.5.2	Geometric Preferences	71
4.5.3	Strategic Behavior	75
4.6	Conclusion and Open Problems	76
4.7	Acknowledgements	77
	Bibliography	79

LIST OF FIGURES

Figure 3.1.	Jealousy Graphs vs. Correlation and Intercorrelation	31
Figure 4.1.	A representation of the reduction from maximum inner product to verifying a stable matching	59
Figure 4.2.	A representation of the reduction from maximum inner product to checking a stable pair	63
Figure 4.3.	A representation of the reduction from maximum inner product to checking a stable pair such that a true maximum inner product instance maps to a false stable pair instance	67

LIST OF TABLES

Table 4.1.	Two-list preferences where no participant receives their top choice in the stable matching	45
Table 4.2.	Two-list preferences where a greedy approach will not work	46
Table 4.3.	Two-list preferences that can be manipulated	48
Table 4.4.	Single-peaked preferences where no participant receives their top choice in the stable matching	70
Table 4.5.	Geometric preferences that can be manipulated	76

ACKNOWLEDGEMENTS

I would like to thank my adviser Mohan Paturi for all of his guidance throughout this graduate program. His insight and mentorship was truly invaluable and I could not have succeeded without it. I would also like to thank my family and friends for their support and encouragement.

Chapter 3, in full, is a reprint of the material as it appears in International Conference on Web and Internet Economics 2013, pages 263–276. Moshe Hoffman, Daniel Moeller, and Ramamohan Paturi. “Jealousy graphs: Structure and complexity of decentralized stable matching.” Springer, 2013. [42] The dissertation author was the primary investigator and author of this paper.

Chapter 4, in part, is a reprint of the material as it appears in International Computer Science Symposium in Russia 2016, pages 298–308. Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. “Subquadratic Algorithms for Succinct Stable Matching.” Springer, 2016. [68] The dissertation author was the primary investigator and author of this paper.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Künnemann, Marvin, Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. “Subquadratic Algorithms for Succinct Stable Matching.” arXiv preprint arXiv:1510.06452 (2016). [59] The dissertation author was the primary investigator and author of this material.

VITA

- 2010 Bachelor of Science, University of Notre Dame
- 2013 Master of Science, University of California, San Diego
- 2016 Doctor of Philosophy, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Exploiting Structure in the Stable Matching Problem

by

Daniel Paul Moeller

Doctor of Philosophy in Computer Science

University of California, San Diego, 2016

Professor Ramamohan Paturi, Chair

Stable matching is a widely studied problem in social choice theory. For the basic centralized case, an optimal quadratic time algorithm is known. However, we present several notions of structure and use them to provide tighter convergence bounds and faster stable matching algorithms for structured instances.

First, we consider the decentralized case, where several natural randomized algorithmic models for this setting have been proposed that have worst case exponential time in expectation. We describe a novel structure associated with a stable matching on a matching market. Using this structure, we are able to provide a finer analysis of the

complexity of a subclass of decentralized matching markets.

We then study the centralized stable matching problem when the preference lists are not given explicitly but are represented in a succinct way. We ask whether the problem becomes computationally easier and investigate other implications of this structure. We give subquadratic algorithms for finding a stable matching in special cases of natural succinct representations of the problem, the d -attribute, d -list, geometric, and single-peaked models. We also present algorithms for verifying a stable matching in the same models. We further show that for $d = \omega(\log n)$ both finding and verifying a stable matching in the d -attribute and d -dimensional geometric models requires quadratic time assuming the Strong Exponential Time Hypothesis. These models are therefore as hard as the general case for large enough values of d .

Chapter 1

Introduction

Social Choice Theory is an important field at the intersection of a wide range of areas, primarily Economics, Social Science, Mathematics, and Computer Science. Initial research focused on voting theory, but it has subsequently expanded to include coalition formation, matching, auctions and other applications.

Much of social choice revolves around resolving the preferences of different individuals to achieve a desired social outcome. For instance, the goal in voting is to aggregate the preferences into a consensus list. In coalition formation, we divide the participants into groups that are acceptable to all involved. Much of the prior work assumes that preferences can be arbitrary because this allows for maximal generality and applicability to a wide range of settings. However, this assumption often leads to impossibility results, as is the case with the well known Arrow's Theorem in voting [8]. This theorem states that there is no voting system that satisfies several common-sense criteria and is not a dictatorship.

To bypass some of these impossibility results, we may often assume there are some restrictions on the preferences. Many of the restrictions are motivated by intuitive or empirical observations. For instance, in voting theory it can be assumed that all participants have single-peaked preferences, which means that each participant has one or more ideal choices with the value of other options monotonically decreasing further away

from the ideal [14]. Arrow's impossibility result no longer holds when all preferences are single-peaked.

In this dissertation, we focus on the stable matching problem, which is a rich problem by itself. In this setting, the participants are divided into two groups, men and women, and each participant has a preference list over the members of the opposite group. The goal is to pair the men and women such that no two people would prefer to be with each other than with their partners. Gale and Shapley [33] first presented the stable matching problem in 1962 as a model for two-sided matching markets. They also describe their deferred acceptance algorithm which efficiently finds a stable matching.

From the time of this seminal paper, stable matching has proved to be a fruitful area of research for several reasons. First, stable matching has many applications ranging from matching buyers to sellers in a market, students to public schools, and residents to hospitals. Each of these are important social and economic problems and there are many real-world examples of such markets. One is the National Resident Matching Program (NRMP) [75], which assigns graduating medical students to residency positions in hospitals. Another is the matching of college students to sororities [69]. Stable matching algorithms are currently used to clear real world markets for these applications.

Moreover, stable matching is of interest from a purely theoretical perspective. On the one hand, it is an approachable topic often used as an example in algorithms courses. However, it also goes much deeper. In his book [56], Knuth presents the relationships between stable matching and many other combinatorial problems. The computational hardness of stable matching has also been investigated. In fact, stable matching is a complete problem in the complexity class CC, which is the set of problems log-space reducible to the comparator circuit value problem [67, 84].

Finally, stable matching has a very rich structure. Gusfield and Irving [38] describe many aspects of this structure. This structure is of interest in various fields. For

example, it can be used to determine how many stable matchings exist for a given problem instance and which would be the most likely outcome, both questions economists ask. From the computer science viewpoint, this structure also enables faster algorithms for problems such as describing the set of stable matchings and the transformations from one stable matching to another. Finally, from a mathematical perspective, the structure in stable matching relates to many basic mathematical constructs such as lattices.

In this work, we consider cases where preferences have structure and use this structure to bypass the current hardness results. We also investigate the limits of this approach. We first discuss the decentralized stable matching setting, where there is no central authority to enforce the participants to follow any stable matching protocol. This applies in many real world settings such as most economic and job markets. Previous work has shown that while distributed agents can converge to a stable matching through a random process, this can take an exponential number of steps [79, 3]. This indicates that we might not expect to find stable outcomes in many decentralized markets. However, we define a notion of structure based on jealousy graphs that allows us to guarantee expected polynomial time convergence when the preferences have certain structural properties. We also demonstrate that markets are likely to have these properties when we make some common assumptions on the preference profiles of the participants. Therefore, we can expect decentralized markets to achieve stable matchings when the preferences allow for better convergence guarantees.

We then consider the centralized setting when the preferences have one of several succinct structures. While the classical Gale-Shapley deferred acceptance algorithm provides the optimal quadratic running time [70, 81, 36], these lower bounds do not necessarily hold when the preferences require subquadratic space to represent. We show that in some cases when the preferences can be implicitly represented by attributes, there are strongly subquadratic algorithms for finding a stable matching. Moreover

there are strongly subquadratic algorithms for verifying if a given matching is stable, which requires quadratic time with arbitrary preferences. We also present a subquadratic algorithm for verifying a stable matching when the agents' preferences are common lists and when they are single-peaked. Finally, we demonstrate that there are limits to the advantage this structure can give by showing that this problem requires quadratic time once the number of attributes becomes superlogarithmic assuming the Strong Exponential Time Hypothesis (SETH).

Chapter 2

Background

2.1 Basic Stable Matching Concepts

We start with the basic definitions of matching markets and stable matchings.

Definition 2.1. (S, P) is a matching market if $S = M \cup W$ for some disjoint sets M, W , $|M| = |W|$ and $P = \{\succ_s\}_{s \in S}$ where, for $s \in M$, \succ_s is a total order over $W \cup \{s\}$, and for $s \in W$, \succ_s is a total order over $M \cup \{s\}$.

We say a matching market has size n if $|M| = |W| = n$.

Definition 2.2. A matching on the set S is a function $\mu : S \rightarrow S$ such that $\forall s \in S$, $\mu(\mu(s)) = s$, $s \in M \Rightarrow \mu(s) \in W \cup \{s\}$ and $s \in W \Rightarrow \mu(s) \in M \cup \{s\}$.

We say that a participant $s \in S$ is unmatched by a matching μ if $\mu(s) = s$. We generally assume that all participants prefer to be matched to anyone than to be unmatched, though this restriction can often be relaxed. Observe that μ can be thought of as a collection of pairs (m, w) if we allow self loops (s, s) for unmatched participants.

Definition 2.3. A matching on the set S is a perfect matching if $\mu(s) \neq s$ for all $s \in S$.

Given a matching, a man and a woman who each preferred the other to their partner would cause the matching to be unstable. Therefore any stable matching must have no such pairs. We call such a pair a blocking pair, defined formally here:

Definition 2.4. Let (S, P) be a matching market and μ be any matching on S . A blocking pair for μ in (S, P) is a pair (m, w) such that $m \in M$, $w \in W$, $\mu(m) \neq w$, $w \succ_m \mu(m)$, and $m \succ_w \mu(w)$.

Definition 2.5. Let (S, P) be a matching market. A matching μ on S is a stable matching for (S, P) if it has no blocking pairs in (S, P) .

2.2 Related Work

2.2.1 The Deferred Acceptance Algorithm

Gale's and Shapley's deferred acceptance algorithm [33] works as follows. While there is an unmatched man m , have m propose to his most preferred woman who has not already rejected him. A woman accepts a proposal if she is unmatched or if she prefers the proposing man to her current partner, leaving her current partner unmatched. Otherwise, she rejects the proposal. This process finds a stable matching in time $O(n^2)$.

2.2.2 Complexity of Stable Matching

It turns out the running time of the deferred acceptance algorithm is optimal under reasonable models of computation. This is to be expected since representing all participants' preferences requires quadratic space. However, even if the preferences are already in memory, finding a stable matching requires quadratic time as well. Additionally, the *verification* problem of testing whether a given matching is stable or not and the *stable pair* problems of checking whether a given pair is in any or all stable matchings also require quadratic time [70, 81, 36].

Subramanian proves that the stable matching problem is actually a complete problem in the complexity class CC [84]. This class is the set of problems equivalent to the comparator circuit value problem, which asks whether a given comparator circuit is

true or false on the given boolean inputs [67]. Thus, reducing stable matching to another problem would imply that problem also had a quadratic lower bound. In fact the stable pair problems are also CC-complete [84]. Le, Cook, and Ye further explore and develop the stable matching problem's relationship with CC [60].

2.2.3 Variations on Stable Matching

Since the seminal work by Gale and Shapley, many variants of this problem have been studied. See [35, 38, 37, 48, 49, 52, 56, 78] for examples. One variation also considered by Gale and Shapley, the many-one stable matching problem, is where members on one side of the market can accept multiple partners [33]. It turns out that the NRMP was already using a form of the deferred acceptance algorithm for their hospital-doctor matchmaking, which is of the many-one type [75]. Many of the results for the standard problem readily extend to this version [38].

Another variation allows for indifference in the preference lists. In the presence of ties, there are three different stability notions. For *super* stability, a blocking pair involves two participants who either prefer each other to their partners or are indifferent. A blocking pair for *strong* stability is the same as for super stability except the preference must be strict for at least one member of a blocking pair. *Weak* stability requires both members of a blocking pair to strictly prefer each other to their partners. It is clear that arbitrarily breaking the ties allows the deferred acceptance algorithm to find a weakly stable matching. However, there is not always a strongly stable or super stable matching. On the other hand, they can be found in polynomial time if they do exist [49, 54].

A third variation involves participants who may find certain partners unacceptable. That is, they would rather be unmatched than matched with those potential partners. This setting is called stable matching with incomplete lists since participants need not rank every member of the other set [38]. Stable matchings always exist in this case, and

the deferred acceptance algorithm can be modified slightly to handle incomplete lists. However, some of the participants may be unmatched even if there are an equal number of men and women. Nevertheless, with strict preferences, the set of unmatched participants is the same for any stable matching on that market [35].

In the presence of both ties and incomplete lists, the stable matching problem becomes somewhat more complex. Weakly stable matchings always exist, however, it is no longer the case that all weakly stable matchings have the same size. Moreover, finding the maximum weakly stable matching is NP-complete [62, 53, 63]. This issue does not apply to strong stability and super stability, and if they exist, stable matchings can be found using similar algorithms to the complete lists algorithms [54, 62].

The *stable roommates problem* is a generalization of the stable matching problem that removes the bipartite restriction. Unlike with bipartite stable matching, there need not always exist a stable roommate matching [33]. However Irving discovered an algorithm that produces a stable matching or identifies that none exists in quadratic time [48]. Subramanian also presents an alternative method [84]. Since every stable matching problem can be realized in the stable roommate setting, stable roommates is at least as hard as stable matching and this is optimal [38]. While we focus on stable matching in this thesis, some of the results immediately generalize to the stable roommates problem as well.

Unless otherwise specified, we deal with the one to one matching case with strict, complete preferences in this dissertation.

2.2.4 Strategic Behavior

Strategy is a key consideration in mechanism design. Ideally, we want to incentivize truthful behavior. In fact, preference misrepresentation is found in real world matching markets that do not encourage truthfulness. (See [69] for an example.) However,

it turns out that there is no strategy-proof mechanism that outputs a stable matching given self-reported preferences [74]. Additionally, if a mechanism outputs the men-optimal stable matching, then the optimal strategy for the women is to report truncated lists which forces the women-optimal stable matching to be the only stable matching [75, 34]. Furthermore, if some women misrepresent their preferences, the partners of all women can only improve whereas the partners of the men can only become worse [9]. The women can actually benefit significantly by falsely reporting their preferences in large, uniform markets [27]. On the other hand, truthfulness is a dominant strategy for the men [74, 76], although they can form coalitions where some, but not all, men in the coalitions can benefit by misrepresentation [43].

Despite these impossibility results, all is not lost. It turns out that under equilibrium, the resulting outcome of the misrepresented preferences will still be stable with regard to the true preferences [61]. Moreover, if the participants have incomplete information about the other participants' preferences, this hinders their ability to gain through misrepresentation [77]. Likewise, if the men have short preference lists, then the probability of the women successfully manipulating the mechanism goes down as the size of the market increases [44, 58]. Also, if the women must rank all men, this limits their ability to gain, but does not completely stop it [85]. Finally, there are also practical obstacles to preference misrepresentation. For instance, [57] and [66] investigate the computational complexity of determining optimal cheating strategies under various constraints.

2.2.5 Structure in Stable Matching

The stable matching problem also has rich structure which often enables efficient algorithms. Most notable is the lattice structure of all stable matchings for a given problem discovered by John Conway and presented in [56] and the accompanying rotation poset

described by Irving and Leather [50]. It turns out that with the man-optimal stable matching as the lower bound and the woman-optimal stable matching as the upper bound, all intermediate stable matchings form points in a lattice structure, where the intermediate stable matchings are preferred by the men to the women-optimal stable matching and by the women to the man-optimal stable matching [56]. Rotations compactly describe the transition from one stable matching to another that is more favorable for the women [50].

The complete rotation poset can be computed in quadratic time and catalyzes fast algorithms for a variety of stable matching problems [37]. For instance, determining if a given pair is a stable pair can be solved in quadratic time. Moreover, all stable pairs can be enumerated in quadratic time [37]. One can also use rotations to efficiently find stable matchings that are more socially optimal, by a variety of measures, than the man-optimal or woman-optimal stable matching. Some examples are the minimum regret stable matching, where the participant who does worst is as well off as possible [37], and the matching with highest average partner ranking [51, 32]. Furthermore, all stable matchings can be enumerated using only linear time per matching and quadratic total space [37]. The utility of this structure is not limited to fast algorithms, as it also serves to provide a #P lower bound for the problem of counting the number of stable matchings [50].

Other notions of structure have been used to distinguish between simple and complex instances of the problem. One commonly considered notion is correlated preferences. In this case, all of the men share the same ranking of the women and vice versa. If a matching market has correlated preferences it is trivial to find the unique stable matching, without using the deferred acceptance algorithm. Of course, most markets will not have perfectly correlated preferences, so several measures of the degree of correlation in preferences have been proposed and exploited for various purposes. See [21],[16], and [28] for examples. We expand further on correlated and the related intercorrelated

preferences in Section 3.4.

In the chapters that follow, we explore other types of structure and demonstrate how these also provide insight into stable matching problems. First we show how one measure of preference structure can provide tighter convergence bounds for a randomized decentralized process. Then we bypass the standard lower bounds for stable matching by investigating instances that have succinctly represented preferences with some inherent structural properties. This leads us to develop algorithms with improved performance as well as several new lower bound results.

Chapter 3

Jealousy Graphs

Most prior stable matching work involves centralized algorithms where the entire set of preferences is known to some central authority. In some cases the algorithms are not totally centralized, but the participants are subject to strict protocols where only one side of the market can make proposals. Nevertheless, many applications of stable matching have no central authority or enforcement of protocols, such as college admissions and the computer scientist job market. Therefore we investigate this problem in a decentralized setting, where members of both sides of the market can make proposals.

One major open question in decentralized stable matching concerns whether natural and efficient algorithms exist. To this end, Yariv argues that natural distributed processes will find stable matchings and provides experimental support [31]. Roth and Vande Vate propose a class of randomized algorithms to model the decentralized setting and show that algorithms in this class converge to a stable matching with probability one [79]. At each step these algorithms match two participants who form a blocking pair (who prefer to be matched with each other over their partners) of the current matching. However, they present no expected time complexity. Ackermann et al. investigate one particular algorithm in this class, the better response algorithm (or random better response dynamics). In each step of this algorithm, one blocking pair is chosen uniformly at random. For this algorithm, they show worst case instances that take exponential time

to reach a stable matching in expectation [3].

Since the better response algorithm is natural but takes exponential time in the worst case, can we find a natural subclass of matching markets which do not require exponential time? Ackermann et al. show that the better response algorithm only requires polynomial time for one class of problem instances, those with correlated preferences [3]. Here, correlated preferences require that a participant obtains the same benefit from a partnership as its partner. This significantly limits the preference structures allowed in the matching market. Therefore, we investigate other structural properties of stable matching markets which facilitate faster convergence.

In this chapter we make progress toward answering the previous question by expanding the subclass of markets with polynomial time convergence guarantees. For this purpose we associate a directed graph, called the *jealousy graph*, with each stable matching. It turns out that this structure is a key factor in determining the convergence time of the better response algorithm. The jealousy graph is a directed graph where a vertex v corresponds to a pair in the stable matching and an edge (u, v) is present if one member of the pair v prefers a member of the pair u to its partner in the stable matching. The strongly connected component graph of this jealousy graph provides a *decomposition* for that stable matching. Our intent is to formalize a notion of structure using jealousy graphs and the corresponding decompositions. In particular, we find that the strongly connected components of this graph give insight into the complexity of that market. Gusfield and Irving provide a structural property of stable matchings which describes the set of stable matchings and the relation between them [38], whereas our structures relate to individual stable matchings and the distributed process by which these stable matchings are achieved.

With a decomposition, we associate a size and depth. Our main result, Theorem 3.2, states that for a matching market of size n with a decomposition of size c and

depth d , the convergence time is $O(c^{O(cd)}n^{O(c+d)})$. Therefore, for constant size and depth decompositions, we demonstrate that the better response algorithm requires only polynomial time in expectation to converge for an expanded class of matching markets. This indicates that the jealousy graph and decomposition structures partially answer the convergence questions of the decentralized stable matching problem. As an application of our work, we demonstrate how Theorem 3.2 provides theoretical justification for the simulated results of Boudreau [16]. We also conjecture that these structures provide a means of predicting which stable matchings are likely to be achieved when there are multiple stable matchings, a question that others in the literature have investigated [17, 18, 31, 71, 13].

3.1 Preliminaries

3.1.1 Better Response Algorithm

The class of algorithms introduced by Roth and Vande Vate [79] involve randomly choosing a blocking pair of the current matching and creating a new matching by matching the participants in the blocking pair with each other. This resolves the chosen blocking pair.

Definition 3.1. A blocking pair (x, y) in a matching μ is resolved by forming a new matching μ' where $\mu'(x) = y$, $\mu'(\mu(x)) = \mu(x)$ if $\mu(x) \neq x$, $\mu'(\mu(y)) = \mu(y)$ if $\mu(y) \neq y$, and $\mu'(s) = \mu(s)$ for $s \notin \{x, y, \mu(x), \mu(y)\}$.

This process is repeated until a stable matching is reached. The *better response algorithm* defined in [3], is the algorithm in this class where the blocking pair is chosen uniformly at random from all blocking pairs of the current matching. Note that this algorithm results in a sequence of matchings. A valid sequence of matchings is any sequence where each matching is formed by resolving one blocking pair in the previous

matching.

We focus on the better response algorithm since the uniform distribution on blocking pairs facilitates our analysis and we believe it provides insight into the more general class of algorithms. This algorithm also serves as a model of a distributed stable matching market. As such, it has been used in simulations of matching market dynamics [16, 17] as well as in modeling locally stable matching dynamics [40, 41].

3.1.2 Jealousy Graph and Related Definitions

The following three concepts are useful since we deal with subsets of matching markets in this chapter.

Definition 3.2. A balanced subset of a matching market (S, P) , $S = M \cup W$, is a subset $S' \subseteq S$ such that $|S' \cap M| = |S' \cap W|$.

Definition 3.3. A matching μ is locally perfect on a balanced subset $S' \subseteq S$ if $\mu(S') \subseteq (S')$ and $\mu \upharpoonright_{S'}$ is a perfect matching on S' .

Definition 3.4. Let μ be a stable matching on a matching market (S, P) . A matching μ' is μ -stable on a balanced subset $S' \subseteq S$ if $\mu' \upharpoonright_{S'} = \mu \upharpoonright_{S'}$.

In order to analyze matching markets, we represent the preference structure as a directed graph. While we lose some of the preference information, we retain critical relationships relative to the stable partners. In section 3 we provide bounds on convergence based on this simpler structure. This graph has also been used to determine reachability if participants enter the market one at a time [25] and is related to the notion of envy graphs in the housing allocation problem [2].

Definition 3.5. The jealousy graph of a stable matching μ on a matching market (S, P) is defined as the graph $J_\mu = (V, E)$ where, for each pair $\{x, \mu(x)\}$, $x \in S$, there is a vertex $v_{\{x, \mu(x)\}} \in V$ and $E = \{(u_{\{x, y\}}, v_{\{x', y'\}}) \mid u_{\{x, y\}}, v_{\{x', y'\}} \in V, \text{ and either } x \succ_{y'} x' \text{ or } y \succ_{x'} y'\}$.

The jealousy graph can provide insight into the complexity of stabilization. For example, suppose the jealousy graph for a stable matching μ is one large clique. Even when all but one pair of the participants are matched with their partner in μ , there are still many blocking pairs. Therefore, the better response algorithm would be unlikely to choose the blocking pair that would result in a stable matching. This greatly hinders convergence to the stable matching.

On the other hand, suppose the jealousy graph for μ is a DAG. Then there is at least one vertex with no incoming edges. This means each partner in the corresponding pair is the other's first preference. Consequently, this will remain a blocking pair until it is resolved, so we would expect such a pair to be resolved in $O(n^2)$ time under the better response dynamics. Moreover, once resolved, the match will remain unbroken since neither partner will ever be involved in any blocking pairs. Ignoring this pair will result in at least one other source vertex of the graph. Inductively, these pairs will be resolved in $O(n^2)$ expected time. This results in an expected convergence time of $O(n^3)$ for the matching market. It should be noted that the class of correlated markets, for which Ackermann et al. prove the better response algorithm requires only polynomial time, falls into this special case.

When it is a DAG, the jealousy graph provides an order in which the pairs will likely be resolved to reach μ , namely, a topological sorted order. However, a matching market might not fall into this extreme case as there could be cycles in the jealousy graph. Therefore, we define a decomposition which is a DAG obtained from the jealousy graph.

Definition 3.6. *Let J_μ be the jealousy graph of a stable matching μ for a matching market (S, P) . A μ -decomposition, ρ_μ is a graph of components of J_μ such that if u, v are in the same strongly connected component of J_μ then they are in the same component in ρ_μ and if edge (A, B) is in ρ_μ then there is a path from a vertex in A to a vertex in B in J_μ .*

We call the strongly connected components of J_μ stable components. Observe that ρ_μ is a directed acyclic graph. Therefore it induces a partial order on the stable components. Sometimes it is simpler to refer to the decomposition as $\rho_\mu = (\Pi, \preceq)$ where Π is a partition of S into sets corresponding to the stable components of ρ_μ and \preceq is the induced partial order on those components. As a slight abuse of notation, we use the term stable component to refer to both the connected component in the decomposition and the set of participants corresponding to this component.

In dealing with partial orders we use the concept of a downset. A *downset* of a partially ordered set Π with partial order \preceq is any set such that for $A, B \in \Pi$, if A is in the set and $B \preceq A$, then B is in the set. The downset of an element $A \in \Pi$ is $Down(A) = \{B \mid B \preceq A\}$. When the elements of Π are sets themselves, as in the case of decompositions, we denote the union of sets in $Down(A)$ as $\mathbf{D}(A) = \bigcup_{B \in Down(A)} B$.

For our complexity results we need the following two notions:

Definition 3.7. *The depth of a stable component A of a μ -decomposition, ρ_μ , is the length of the longest path in ρ_μ from any source vertex to v_A . The depth of ρ_μ is defined as $\max_{A \in \rho_\mu} depth(A)$.*

We say that a stable component A is on level j if $depth(A) = j$. Minimal stable components are on level 0. Intuitively, we would expect components on lower levels to converge to the stable matching sooner than those on higher levels.

Definition 3.8. *The size of a μ -decomposition, ρ_μ , is defined as $\max_{A \in \rho_\mu} size(A)$.*

Intuitively, components with smaller sizes can have less internal thrashing so they will converge to the stable matching more quickly than larger components.

3.2 Structural Results

3.2.1 Any Digraph can be a Jealousy Graph

These structural notions would not be very enlightening if all matching markets had similar jealousy graphs and decompositions. However, the following result shows that any directed graph is the jealousy graph associated with a stable matching for some matching market.

Theorem 3.1. *Given any directed graph G with n vertices, there is a set $S = \{m_i, w_i \mid i = 1, 2, \dots, n\}$ and preferences $P = \{\succ_{m_i}, \succ_{w_i} \mid 1 \leq i \leq n\}$ such that (S, P) is a matching market with a stable matching μ where $\mu(m_i) = w_i$ and $J_\mu = G$.*

Proof. Let $G = (V, E)$ and $S = \{m_i, w_i \mid i = 1, 2, \dots, n\}$. Arbitrarily index the vertices in V as v_1, v_2, \dots, v_n . We design our preferences such that if $\mu(m_i) = w_i$ for $i = 1, 2, \dots, n$, μ is a stable matching and v_i is the vertex corresponding to $\{m_i, w_i\}$. Now define $P = \{\succ_{m_i}, \succ_{w_i} \mid 1 \leq i \leq n\}$ as follows. First, for every woman w_i , define \succ_{w_i} such that $m_i \succ_{w_i} m_j$ for $j \neq i$. The remaining ordering can be arbitrary. For every man, define \succ_{m_i} such that $w_j \succ_{m_i} w_i \Leftrightarrow (v_j, v_i) \in E$. The ordering among elements within $\{w_j \mid w_j \succ_{m_i} w_i\}$ and $\{w_j \mid w_i \succ_{m_i} w_j\}$ can be arbitrary.

To see that μ is indeed a stable matching, observe that under μ all women are matched with their top choice. Therefore, no woman has incentive to deviate, so there can be no blocking pairs and μ is a stable matching on S, P .

All that remains is to show $J_\mu = G$. Now in $J_\mu = (V', E')$ let the vertices be denoted $V' = v'_1, v'_2, \dots, v'_n$. We let v'_i correspond to $\{m_i, w_i\}$ for $i = 1, 2, \dots, n$. Since all women are matched with their top preference by μ , the women are not responsible for any edges in J_μ . Therefore, $(v'_i, v'_j) \in E' \Leftrightarrow w_i \succ_{m_j} w_j \Leftrightarrow (v_i, v_j) \in E$ by the way we defined \succ_{m_j} for all i, j . Thus if we let $v_i = v'_i$ we have equivalent graphs. \square

3.2.2 Properties of Decompositions

In this section we prove several structural properties of the jealousy graphs and decompositions essential to our main convergence result. The first property says that if there is a path from one vertex to another in the jealousy graph, then the first vertex must be in the downset of any component containing the second vertex.

Lemma 3.1. *Given a matching market (S, P) with a stable matching μ , let J_μ be the jealousy graph associated with μ . Let $v_{\{m, w\}}$ and $v_{\{m', w'\}}$ be vertices in J_μ . Suppose $v_{\{m', w'\}} \in A$ for a stable component A of a μ -decomposition $\rho_\mu = (\Pi, \preceq)$. If there is a path from $v_{\{m, w\}}$ to $v_{\{m', w'\}}$, then $m, w \in \mathbf{D}(A)$.*

Proof. Let v_1, v_2, \dots, v_k be the vertices along the path such that $v_1 = v_{\{m, w\}}$ and $v_k = v_{\{m', w'\}}$. Let A_i be the stable component containing v_i . Then since the partial order is induced by the edges between components of J_μ , either $A_i = A_{i+1}$ or $A_i \preceq A_{i+1}$. Therefore, by transitivity $A_1 \preceq A_k = A$. Thus $m, w \in \mathbf{D}(A)$ because $v_{\{m, w\}} \in A_1$. \square

Using this lemma, we prove that no member of a stable component can prefer anyone outside of the downset of that component to his stable partner.

Lemma 3.2. *Given a matching market (S, P) with a stable matching μ , let $\rho_\mu = (\Pi, \preceq)$ be a μ -decomposition. For $A \in \Pi$, $a \in A$, $s \in S - \mathbf{D}(A)$, $\mu(a) \succ_a s$.*

Proof. Let $A \in \Pi$ and $a \in A$. Suppose there is some $s \in S - \mathbf{D}(A)$ such that $s \succ_a \mu(a)$. Then since $s \notin A$, there are distinct vertices in J_μ , v_a, v_s corresponding to the pair with a and the pair with s , respectively. Edge (v_s, v_a) must also be in J_μ since $s \succ_a \mu(a)$. Thus there is a path in J_μ from v_s to v_a , so by Lemma 3.1, $s \in \mathbf{D}(A)$. This is a contradiction. \square

A further property is that if there are two stable matchings with distinct decompositions, the intersection of the downsets of stable components must be mapped to itself in both stable matchings.

Lemma 3.3. *Given a matching market (S, P) with stable matchings μ, μ' , let ρ_μ and $\rho_{\mu'}$ be respective decompositions. Let A be $\mathbf{D}(X)$ for some stable component X of ρ_μ and B be $\mathbf{D}(Y)$ for some stable component Y of $\rho_{\mu'}$. Then $\mu(A \cap B) = \mu'(A \cap B) = A \cap B$.*

Proof. Suppose there is $x \in A \cap B$ but $\mu(x) \in A - B$. By Lemma 3.2, $\mu'(x) \succ_x \mu(x)$. In that case, $\mu'(x) \in A \cap B$, also by Lemma 3.2. Now since x prefers $\mu'(x)$ to $\mu(x)$, $\mu(\mu'(x)) \succ_{\mu'(x)} x$ or else $(x, \mu'(x))$ forms a blocking pair for μ . Again by Lemma 3.2, $\mu(\mu'(x)) \in A \cap B$. Continuing in this manner gives an infinite sequence of participants $x, \mu'(x), \mu(\mu'(x)), \mu'(\mu(\mu'(x))), \dots \in A \cap B$. These are all distinct since $\mu(x) \neq \mu'(x)$ and both μ and μ' are bijective, which is a contradiction since $A \cap B$ is finite. Therefore $\mu(A \cap B) = \mu'(A \cap B) = A \cap B$. \square

Our final result shows that forming a stable matching on the downset of a stable component cannot increase the size or depth of the decomposition of another stable matching.

Lemma 3.4. *Given a matching market (S, P) with stable matchings μ, μ' , let ρ_μ and $\rho_{\mu'}$ be respective decompositions. Suppose the size of ρ_μ is c and the depth is d . Let A be a stable component of $\rho_{\mu'}$. Then there is a stable matching μ'' such that $\mu'' \upharpoonright_{\mathbf{D}_{\mu'}(A)} = \mu' \upharpoonright_{\mathbf{D}_{\mu'}(A)}$ and $\mu'' \upharpoonright_{S - \mathbf{D}_{\mu'}(A)} = \mu \upharpoonright_{S - \mathbf{D}_{\mu'}(A)}$. There is also a μ'' -decomposition on $S - \mathbf{D}_{\mu'}(A)$ of size at most c and depth at most d .*

Proof. Let μ'' be such that $\mu'' \upharpoonright_{\mathbf{D}(A)} = \mu' \upharpoonright_{\mathbf{D}(A)}$ and $\mu'' \upharpoonright_{S - \mathbf{D}(A)} = \mu \upharpoonright_{S - \mathbf{D}(A)}$. Clearly there are no blocking pairs involving two members of $\mathbf{D}(A)$ or else μ' would not be stable and there are no blocking pairs between two members of $S - \mathbf{D}(A)$ or else μ would not be stable. Finally, by Lemma 3.2 no member of $\mathbf{D}(A)$ can prefer any member of $S - \mathbf{D}(A)$ to his partner in μ' . Therefore there can be no blocking pairs between a member of $\mathbf{D}(A)$ and a member of $S - \mathbf{D}(A)$ so μ'' is indeed a stable matching.

By Lemma 3.3, for each stable component B of ρ_μ , $\mu(B - \mathbf{D}(A)) = B - \mathbf{D}(A)$. The set $\{B - \mathbf{D}(A) \mid B \in \rho_\mu\}$ forms a partition of $S - \mathbf{D}(A)$ and, paired with the same partial order as ρ_μ , forms a decomposition. Clearly the size has not increased since the sets in the partition are no larger and the depth has not increased since the decomposition has the same partial order as ρ_μ . \square

3.3 Convergence

In this section we prove our convergence result. The proof uses two main ideas. First, in the following sequence of lemmas, we show that a stable component will converge to a locally perfect matching in time that is only polynomially dependent on the size of the entire market. Then the proof of Theorem 3.2 uses this to bound the time it takes for all components of the decomposition to reach a stable matching.

For this section we assume (S, P) is a matching market of size n , μ is a stable matching on S , and (Π, \preceq) be a μ -decomposition.

The following lemma says that if a matching is not locally perfect on a stable component of a μ -decomposition, then there is a blocking pair which is in μ between two members of that component.

Lemma 3.5. *Let $A \in \Pi$ and $X = \mathbf{D}(A) - A$. Let μ' be the current matching. If μ' has no matches between members of X and members of A and μ' is not locally perfect on A , then there is a blocking pair (x, y) for μ' such that $x, y \in A$ and $\mu(x) = y$.*

Proof. Since μ' is not a locally perfect matching on A there must be some $x_0 \in A$ such that $\mu'(x_0) = x_0$ or $\mu'(x_0) \in S - X - A$. Let $y_0 = \mu(x_0)$. Now since μ is a stable matching, $y_0 \succ_{x_0} \mu'(x_0)$. If $x_0 \succ_{y_0} \mu'(y_0)$ then (x_0, y_0) is a blocking pair of μ' and $\mu(x_0) = y_0$.

Otherwise $\mu'(y_0) \succ_{y_0} x_0$, so $\mu'(y_0) \in \mathbf{D}(A)$. In fact, $\mu' \in A$ since μ' has no matches between members of A and X . Let $x_1 = \mu'(y_0)$ and $y_1 = \mu(x_1)$. Since μ is a

stable matching, $y_1 \succ_{x_1} y_0$ or else (x_1, y_0) would form a blocking pair for μ . Now if $x_1 \succ_{y_1} \mu'(y_1)$, (x_1, y_1) is a blocking pair of μ' and $\mu(x_1) = y_1$ so we have our result. Otherwise we repeat in the same manner to form a sequence of pairs $\{(x_i, y_i)\}$ such that $x_i, y_i \in A$, $\mu(x_i) = y_i$, $\mu'(y_i) = x_{i+1}$, $y_i \succ_{x_i} \mu'(x_i)$, and $x_{i+1} \succ_{y_i} x_i$ for all i . But this cannot cycle since no participant is repeated. This is because at each step we add a new pair x_i, y_i where $\mu(x_i) = y_i$ and either $\mu'(x_0) = x_0$ or $\mu'(x_0) \notin A$, so x_0 cannot be repeated. Furthermore, it cannot go forever since A is finite. Therefore the sequence must terminate at some index k and (x_k, y_k) is a blocking pair for μ' . \square

Next we place a lower bound on the probability that we make some progress toward the μ -stable matching when a stable component of the decomposition is not in a locally perfect matching.

Lemma 3.6. *Let $A \in \Pi$ be a stable component of size at most c and $X = \mathbf{D}(A) - A$. Let μ' be any matching on S that is not a locally perfect matching on A . Then starting from μ' , if no matches are formed between a member of A and a member of X , the probability that the first blocking pair resolved between two members of A is a pair in μ is at least $\frac{1}{c^2}$.*

Proof. Lemma 3.5 shows there will be one blocking pair which is in μ until the matching becomes locally perfect on A . In order for the matching to become locally perfect on A , a blocking pair must be resolved between two members of A . Therefore since there will be at most c^2 blocking pairs involving two members of A and at least one of them is in μ , there is a $\frac{1}{c^2}$ probability that the first blocking pair resolved between members of A is in μ . \square

Using this lemma, we bound the probability that a component of the decomposition will make some progress toward the μ -stable matching each time the matching is not locally perfect on it.

Lemma 3.7. *Let $A \in \Pi$ be a stable component of size at most c and $X = \mathbf{D}(A) - A$. Let μ_0 be any matching on S such that $\mu_0 \upharpoonright_A$ contains m of the pairs in μ where $0 \leq m < c$. Let $\mu_0, \mu_1, \dots, \mu_t$ be any valid sequence of matchings under the better response dynamics starting from μ_0 such that*

1. μ_t is locally perfect on A
2. μ_i is not locally perfect on A for some $i, 0 \leq i < t$
3. μ_k does not have any matches between a member of A and a member of X for some $k, 0 \leq k \leq t$

Then the probability that $\exists j, 0 < j \leq t, \mu_j \upharpoonright_A$ contains at least $m + 1$ of the pairs in μ is at least $\frac{1}{c^A}$.

Proof. Assume $\mu_0, \mu_1, \dots, \mu_t$ is such a sequence, and i is the first index such that μ_i is not locally perfect. Without loss of generality assume $k = t$ is the first index $k > i$ such that μ_k is locally perfect on A . This assumption is valid because, if there is at least a probability p of some event occurring in a subsequence, then there is clearly at least a probability p of that event occurring in the entire sequence.

There are two cases: either μ_0 is locally perfect on A or not.

case i: Assume μ_0 is not locally perfect, so $i = 0$. Then in order to reach μ_t there must be at least one match formed between two members of A . Let $j > 0$ be the first index in the sequence such that μ_j was formed by resolving a blocking pair between two members of A . Since no one in A prefers anyone in $S - \mathbf{D}(A)$ to his partner in μ , $\mu_{j-1} \upharpoonright_A$ has m pairs in μ . By lemma 3.6 there is at least $\frac{1}{c^2}$ probability that the first blocking pair resolved between two members of A is in μ . This will result in $\mu_j \upharpoonright_A$ having $m + 1$ pairs in μ .

case ii: If μ_0 is locally perfect, so $i > 0$. There are two ways to transition from μ_{i-1} to μ_i . One is for a blocking pair of μ_{i-1} between a member of A and a member

of $S - X - A$ to be resolved. Since this cannot involve a member of A who is with his partner in μ according to μ' , $\mu_i \upharpoonright_A$ has m pairs that are in μ . Therefore this case reduces to the first case where the initial matching is not perfect.

The other way to transition from μ_{i-1} to μ_i is for a blocking pair between two members of A to be resolved, leaving two unmatched members of A , say x, y . The blocking pair cannot involve two pairs of μ or else it would be a blocking pair for μ . If it involves no pairs of μ then again this case reduces to the first case.

In the last case, $\mu_i \upharpoonright_A$ has $m - 1$ pairs that are in μ . We cannot reach μ_t without resolving a blocking pair between two members of A . Let $l > i$ be the first index after i in the sequence such that μ_l was formed by resolving a blocking pair between two members of A . Then $\mu_{l-1} \upharpoonright_A$ must have $m - 1$ pairs that are in μ . By lemma 3.6 there is at least $\frac{1}{c^2}$ probability that $\mu_l \upharpoonright_A$ has m pairs that are in μ . If this occurs, the blocking pair resolved to transition to μ_l cannot involve both x and y because they are not partners in μ . Thus, at least one of x or y is still not matched to someone in A . Therefore, μ_l is not a locally perfect matching on A . Then by the first case, we have at least $\frac{1}{c^2}$ probability that for some j , $l < j \leq t$, $\mu_j \upharpoonright_A$ has $m + 1$ pairs that are in μ . This gives us a total probability of at least $\frac{1}{c^4}$ that $\mu_j \upharpoonright_A$ has $m + 1$ pairs that are in μ for some j , $0 < j \leq t$. \square

We now bound the expected number of times each stable component will have to become not locally perfect before it becomes μ -stable.

Lemma 3.8. *Let $A \in \Pi$ be a stable component of size at most c and $X = \mathbf{D}(A) - A$. Let μ' be any matching on S . Then starting from μ' , if no matches are formed between a member of A and a member of X , the expected number of distinct times the matching needs to transition from a locally perfect matching on A to a matching that is not locally perfect on A before it reaches a μ -stable matching on A is at most $c^{4(c+1)}$.*

Proof. Consider a Markov chain with states $\{0, 1, \dots, c\}$ where state i represents a

matching whose restriction to A has i pairs in μ . Let t_i be the expected number of times, starting from state i , that the matching transitions from a locally perfect matching on A to a matching that is not locally perfect on A before it reaches a μ -stable matching on A . Then $t_c = 0$ since state c represents a μ -stable matching. For all other states, by lemma 3.7, we have at least a $\frac{1}{c^4}$ probability of reaching state $i+1$ from state i after one or fewer transitions from a locally perfect matching on A to a matching that is not locally perfect on A . In the worst case, we will move to state 0 after one such transition with the remaining probability. This leads to the formula $t_i \leq \frac{c^4-1}{c^4}t_0 + \frac{1}{c^4}t_{i+1} + 1$ for $i = 0, 1, \dots, n$.

We need to upper bound t_0 since 0 is the farthest state from c . Now $t_0 \leq t_1 + c^4$. Furthermore if $t_0 \leq t_i + \sum_{j=1}^i c^{4j}$, then

$$t_0 \leq \frac{c^4-1}{c^4}t_0 + \frac{1}{c^4}t_{i+1} + 1 + \sum_{j=1}^i c^{4j}$$

so

$$\frac{1}{c^4}t_0 \leq \frac{1}{c^4}t_{i+1} + \sum_{j=1}^i c^{4j} + 1$$

and

$$t_0 \leq t_{i+1} + c^4 \left(\sum_{j=1}^i c^{4j} + 1 \right) = t_{i+1} + \sum_{j=1}^{i+1} c^{4j}$$

Therefore $t_0 \leq t_c + \sum_{j=1}^i c^{4j} = \sum_{j=1}^i c^{4j} < c^{4(c+1)}$. □

The final lemma we need shows that when the matching is not locally perfect on a stable component of the decomposition, it will reach a perfect matching in time that depends only linearly in n in expectation, provided there is no interference from members of lower stable components.

Lemma 3.9. *Let $A \in \Pi$ be a stable component of size at most c and $X = \mathbf{D}(A) - A$. Let*

μ' be any matching on S which is not locally perfect on A . Then starting from μ' , if no matches are formed between a member of A and a member of X , the expected time reach a matching which is locally perfect on A is at most cn^{2c} .

Proof. Lemma 3.5 implies that for any given matching, either the matching is locally perfect on A or there is a blocking pair between two members of A which is a pair in μ . Since the size of A is at most c , there are at most c such pairs. Therefore if all of them are resolved in c consecutive steps, the resulting matching will be locally perfect on A . Alternatively if after fewer than c steps of resolving blocking pairs that are in μ we reach a matching with no such blocking pairs, then the matching must already be locally perfect on A . For any given matching there are at most n^2 total blocking pairs so the probability of resolving a blocking pair between two members of A that is a pair in μ is at least $\frac{1}{n^2}$. But then the probability of resolving up to c of them and reaching a locally perfect matching in c or fewer steps is at least $\frac{1}{n^{2c}}$.

Therefore, in expectation we will have to repeat the process of making c steps at most n^{2c} times before reaching a locally perfect matching on A . This leads to at most cn^{2c} steps in expectation. \square

Finally we show that the expected convergence time for the better response dynamics is linear in the total number of participants but possibly exponential in the size of the largest stable component and depth of the decomposition. The special case where the size of the decomposition is 1 includes the correlated preferences of Ackermann et al.

Theorem 3.2 (Convergence). *Suppose μ is a stable matching. Suppose the depth of (Π, \preceq) is d and the size of the largest stable component of Π is no more than c . Then the expected time to converge to a stable matching is $O(c^{O(cd)}n^{O(c+d)})$. If $c = 1$, then the expected time is $O(n^3)$.*

Proof. Suppose μ' is another stable matching. First, suppose that for any stable component A' of a μ' -decomposition, a μ' -stable matching is never reached on $\mathbf{D}_{\mu'}(A')$.

Consider the μ -decomposition graph for (Π, \preceq) . Recall that a stable component A is on level j if $\text{depth}(A) = j$. For convenience, let level $d + 1$ be an empty dummy level at the top. Since the depth is d , there are exactly $d + 1$ levels. We proceed by bounding the expected time for one level to reach a μ stable matching, and then recurse on the higher levels.

Let $T(l)$ denote the expected time for the participants in stable components on levels l and above to reach a stable matching without resolving blocking pairs involving any members of stable components on lower levels. Let n_l be the number of stable components on level l . Note that since there are at most n stable components of \mathcal{D} , $n_1 + \dots + n_d \leq n$. We show that $T(0) = O(c^{O(cd)} n^{O(c+d)})$.

First observe that $T(d + 1) = 0$ since there are no stable components at level $d + 1$.

Now consider $T(l)$ for $l < d + 1$.

When one of the n_l stable components A on level l is not in a locally perfect matching. Then by Lemma 3.9, we know it will take cn^{2c} steps in expectation to reach a locally perfect matching on A . Also, by lemma 3.7 we know it has at least $\frac{1}{c^2}$ probability of reaching a matching whose restriction to A has a greater number of pairs that are in μ than the current matching, before it reaches a locally perfect matching.

On the other hand, when all n_l stable components are in locally perfect matchings, then there are two cases:

If there is a blocking pair between two members of stable components on level l it will remain there until the matching becomes not locally perfect on at least one stable component on level l . Since there are at most n^2 blocking pairs, it will take at most n^2 steps in expectation for the matching to become not locally perfect on at least one stable

component on level l .

If there are no such blocking pairs, it might be required for the higher levels to reach a stable matching before exposing a blocking pair involving a participant on level l . If no matches are formed involving any members of components on level l or lower, the expected time for the remaining stable components to reach a stable matching is given by $T(l+1)$. Once the higher levels have reached a stable matching, the only blocking pairs not involving members of levels below l are between a member of a stable component on level l and a member of a stable component on a higher level. Unless all stable components on level l and above are in a stable matching, at least one such blocking pair must exist. Therefore it will only take 1 more step to reach a matching which is not locally perfect on one stable component on level l .

Consequently, it will take at most $n^2 + T(l+1) + 1$ steps to reach a matching that is not locally perfect on one stable component on level l . Again, by Lemma 3.9, we know it will take cn^{2c} steps in expectation to reach a locally perfect matching on A . By Lemma 3.8, we know in expectation, for each stable component on level l , it will take at most $c^{4(c+1)}$ transitions from a locally perfect matching to a matching which is not locally perfect on that stable component it reaches a μ -stable matching. This means that in expectation it will take at most $n_l c^{4(c+1)}$ of these transitions total before all stable components on level l reach a μ -stable matching.

Therefore, in the worst case, it will take $(n^2 + T(l+1) + 1)$ steps to transition from a locally perfect matching to a matching that is not locally perfect on one of the stable components on level l . Then it will take at most cn^{2c} steps to reach a matching which is locally perfect on that stable component. Furthermore, this process needs to be repeated no more than $n_l c^{4(c+1)}$ times in expectation in order for all stable components on level l to reach a μ -stable matching.

Once all stable components on level l have reached a μ -stable matching, all that

remains is for the higher levels to reach a stable matching, which takes $T(l+1)$ time in expectation.

This yields the following formula:

$$T(l) \leq n_l c^{4(c+1)} (cn^{2c} + n^2 + T(l+1) + 1) + T(l+1) \leq 2n_l c^{4(c+1)} (cn^{2c} + T(l+1))$$

Solving this recursion for $T(0)$, we obtain

$$\begin{aligned} T(0) &\leq 2n_0 c^{4(c+1)} (cn^{2c} + T(1)) \\ T(0) &\leq (cn^{2c}) \sum_{i=1}^{d+1} (2c^{4(c+1)})^i \prod_{j=0}^{i-1} n_j \end{aligned}$$

so since $n_i + 1 \leq O(n)$ for all i , $T(0) = O(c^{O(cd)} n^{O(c+d)})$.

This is the expected time to reach the stable matching μ . Now suppose for some stable component A' of a μ' -decomposition for some other stable matching μ' , a μ' -stable matching is reached on $\mathbf{D}_{\mu'}(A')$. By Lemma 3.4, this will not increase the size or depth of the remaining decomposition. Therefore, if this happens before μ is reached, it will only decrease the convergence time.

Finally, as a special case assume $c = 1$. In this case a locally perfect matching on a stable component is a μ -stable matching. By lemma 3.9 it will take at most n^2 steps for a stable component on level l to reach a μ -stable matching. Since there are n_l components on level l , $T(l) \leq n_l n^2 + T(l+1) \leq \sum_{i=1}^{d-l} n_i n^2$ so $T(0) \leq \sum_{i=1}^d n_i n^2 = n^3$. \square

3.4 Correlated and Intercorrelated Preferences

We have shown bounds on convergence time but this is only relevant if there is variation in the jealousy graph structures of real markets. While randomly generated

preferences tend to have decompositions that are close to the trivial decomposition, which is the entire set, real-world markets tend to have some structure. Here we show that two classes of preferences found in real world markets, correlated and intercorrelated preferences, exhibit decompositions with small size components. Partially correlated preferences are often used by modelers [16, 23] and are natural in many matching markets (e.g. mate selection) where preferences are based on a mixture of universally desirable features (e.g. intelligence) and idiosyncratic tastes (e.g. shared hobbies). Note that the correlated preferences discussed here differ from the correlated preferences of Ackermann et al. Intercorrelation exists when the preferences of the men relate to the preferences of the women. See [19] for examples of markets with intercorrelation. Boudreau showed that more correlation and intercorrelation lead to faster convergence of the better response algorithm [16]. We provide similar plots in Figures 3.1b and 3.1d. Theorem 3.2 provides theoretical justification for these simulated results.

As described in [21, 23], correlated preferences are generated using scores of the form:

$$S_{mw} = \eta_{mw} + UI_w$$

where S_{mw} is the score man m gives woman w composed of his individual score η_{mw} and a correlation factor $U \in [0, \infty)$ multiplied by the consensus score of w , I_w . η_{mw} and I_w are chosen uniformly at random from $[0, 1]$. The men then rank the women in order from lowest score to highest. Women's preferences are generated analogously. For various values of U we generate 100 preferences with correlation factor U . For each set of preferences we find the decomposition with smallest size and report the average of these sizes. We also compute the average minimal depth in the same manner. The results are shown in figure 3.1a. At $U = 0$, the average size is close to n and the depth is close to 1. As U goes to ∞ , the average size approaches 1 and the depth approaches n . These

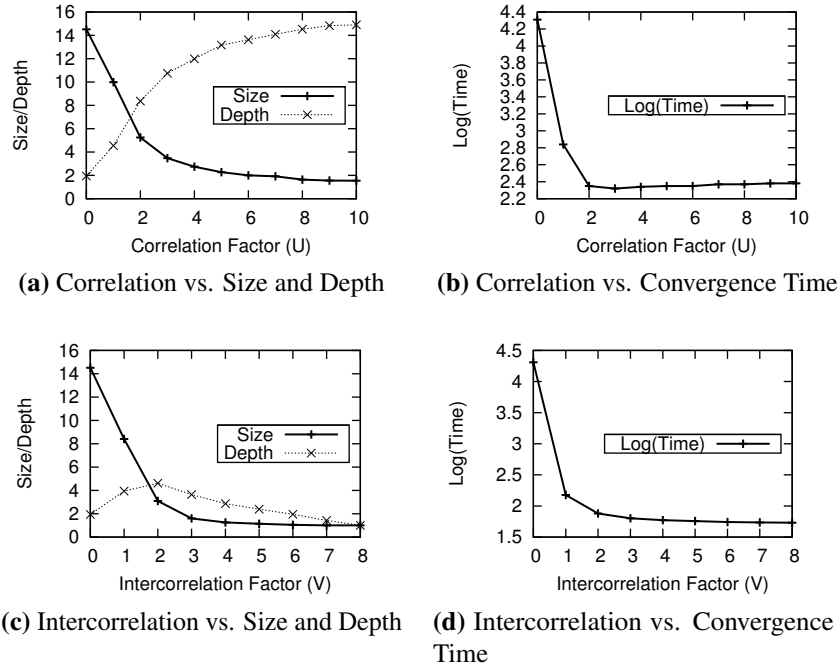


Figure 3.1. Jealousy Graphs vs. Correlation and Intercorrelation: (a) The jealousy graph parameters change as preferences become more correlated. (b) Convergence time decreases as preferences become more correlated. (c) The jealousy graph parameters change as preferences become more intercorrelated. (d) Convergence time decreases as preferences become more intercorrelated.

are the parameters of perfectly correlated preferences. This shows that as the amount of correlation varies, so do the size and depth of the decompositions. Figure 3.1b shows the log of the average convergence time over 100 trials for each of the 100 correlated preferences generated.

As in [19], intercorrelated preferences can be generated using scores of the form:

$$S_{m_i w_j} = \eta_{m_i w_j} + V * |i - j|_n$$

where $S_{m_i w_j}$ is the score man m_i gives woman w_j . As with correlated preferences, $\eta_{m_i w_j}$ is his individual score. Here V is the intercorrelation factor and $|i - j|_n = \min(|i - j|, n - |i - j|) / (\frac{n}{2})$ represents the “distance” man m_i is from woman w_j . 3.1c and 3.1d

are generated in the same manner as 3.1*a* and 3.1*b*, respectively. These plots show that as preferences become more intercorrelated, the size and depth of the decompositions decrease. As Theorem 3.2 explains, this decreases the convergence time of the better response algorithm as intercorrelation increases.

3.5 Conclusion and Open Problems

We have introduced a new way of viewing stable matching problems in terms of their jealousy graphs and μ -decompositions. We demonstrate that these concepts are useful in analyzing the convergence time of the better response algorithm and guarantee polynomial convergence on a subclass of matching markets. Furthermore, these theoretical results apply to a broad range of markets since they provide a notion of structure which extends beyond the well-studied notions of correlation and intercorrelation.

One open question involves the exponential dependency on the depth of the decomposition. While we know that the exponential dependency on size cannot be removed, it remains an open question whether we can improve this bound in terms of the depth. Another open problem concerns which matching is most likely to be reached. Since our result provides a method of classifying the expected convergence time of the better response algorithms in terms of the decompositions of the stable matchings, we conjecture that matchings with decompositions that have small size and depth are more likely to be reached than ones with large size and depth. Finally, we could explore the decentralized strategic implications, as in [80], when restricting the preferences to have jealousy graphs and decompositions with small size and depth.

3.6 Acknowledgements

Chapter 3, in full, is a reprint of the material as it appears in International Conference on Web and Internet Economics 2013, pages 263–276. Moshe Hoffman,

Daniel Moeller, and Ramamohan Paturi. “Jealousy graphs: Structure and complexity of decentralized stable matching.” Springer, 2013. [42] The dissertation author was the primary investigator and author of this paper.

Chapter 4

Succinct Preference Models

For arbitrary preferences, the deferred acceptance algorithm of Gale and Shapley is optimal and even verifying that a given matching is stable requires quadratic time [70, 81, 36]. However, in many applications the preferences are not arbitrary and can have more structure. For example, top doctors are likely to be universally desired by residency programs and students typically seek highly ranked schools. In these cases participants can represent their preferences succinctly. It is natural to ask whether the same quadratic time bounds apply with compact and structured preference models that have subquadratic representations. This will provide a more nuanced understanding of where the complexity lies: Is stable matching inherently complex, or is the complexity merely a result of the large variety of possible preferences? To this end, we examine several restricted preference models with a particular focus on two originally proposed by Bhatnagar et al. [12], the d -attribute and d -list models. Using a wide range of techniques we provide algorithms and conditional hardness results for several settings of these models.

In the d -attribute model, we assume that there are d different attributes (e.g. income, height, sense of humor, etc.) with a fixed, possibly objective, ranking of the men for each attribute. Each woman's preference list is based on a linear combination of the attributes of the men, where each woman can have different weights for each

attribute. Some women may care more about, say, height whereas others care more about sense of humor. Men’s preferences are defined analogously. This model is applicable in large settings, such as online dating systems, where participants lack the resources to form an opinion of every other participant. Instead the system can rank the members of each gender according to the d attributes and each participant simply needs to provide personalized weights for the attributes. The combination of attribute values and weights implicitly represents the entire preference matrix. Bogomolnaia and Laslier [15] show that representing all possible $n \times n$ preference matrices requires $n - 1$ attributes. Therefore it is reasonable to expect that when $d \ll n - 1$, we could beat the worst case quadratic lower bounds for the general stable matching problem.

In the d -list model, we assume that there are d different rankings of the men. Each women selects one of the d lists as her preference list. Similarly, each man chooses one of d lists of women as his preference list. This model captures the setting where members of one group (i.e. student athletes, sorority members, engineering majors) may all have identical preference lists. Mathematically, this model is actually a special case of the d -attribute model where each participant places a positive weight on exactly one attribute. However, its motivation is distinct and we can achieve improved results for this model.

Chebolu et al. prove that approximately counting stable matchings in the d -attribute model for $d \geq 3$ is as hard as the general case [24]. Bhatnagar et al. showed that sampling stable matchings using random walks can take exponential time even for a small number of attributes or lists but left it as an open question whether subquadratic algorithms exist for these models [12].

We show that faster algorithms exist for finding a stable matching in some special cases of these models. In particular, we provide subquadratic algorithms for the d -attribute model, where all values and weights are from a small set, and the one-sided

d -attribute model, where one side of the market has only one attribute. These results show we can achieve meaningful improvement over the general setting for some restricted preferences.

While we only provide subquadratic algorithms to find stable matchings in special cases of the attribute model, we have stronger results concerning verification of stable matchings. We demonstrate optimal subquadratic stability testing algorithms for the d -list and boolean d -attribute settings as well as a subquadratic algorithm for the general d -attribute model with constant d . These algorithms provide a clear distinction between the attribute model and the general setting. Moreover, these results raise the question of whether verifying and finding a stable matching are equally hard problems for these restricted models, as both require quadratic time in the general case.

Additionally, we show that the stable matching problem in the d -attribute model for $d = \omega(\log n)$ cannot be solved in subquadratic time under the Strong Exponential Time Hypothesis (SETH) [45, 47]. We show SETH-hardness for both finding and verifying a stable matching and for checking if a given pair is in any or all stable matchings, even when the weights and attributes are boolean. This adds the stable matching problem to a growing list of SETH-hard problems, including Fréchet distance [20], edit distance [10], string matching [1], k -dominating set [72], orthogonal vectors [86], and vector domination [46]. Thus the quadratic time hardness of the stable matching problem in the general case extends to the more restricted and succinct d -attribute model. This limits the space of models where we can hope to find subquadratic algorithms.

We further present several results in related succinct preference areas. Single-peaked preferences are commonly used to model preferences in social choice theory because of their simplicity and because they often approximate reality. Essentially, single-peaked preferences require that everyone agree on a common spectrum along which all alternatives can be ranked. However, each individual may have a different ideal

choice and prefers the “closest” alternatives. A typical example is the political spectrum where candidates fall somewhere between liberal and conservative. In this setting, voters tend to prefer the candidates that are closer to their own ideals. As explained below, these preferences can be succinctly represented. Bartholdi and Trick [11] present a subquadratic time algorithm for stable roommates (and stable matching) with narcissistic, single-peaked preferences. In the narcissistic case, the participants are located at their own ideals. This makes sense in some applications but is not always realistic. We provide a subquadratic algorithm to verify if a given matching is stable in the general single-peaked preference model. Chung uses a slightly different model of single-peaked preferences where a stable roommate matching always exists [26]. In this model the participants would rather be unmatched than matched with someone further away from their ideal than they are themselves, leading to incomplete preference lists.

We extend our algorithms and lower bounds for the attribute model to the geometric model where preference orders are formed according to euclidean distances among a set of points in multi-dimensional space. Arkin et al. [7] derive a subquadratic algorithm for stable roommates with narcissistic geometric preferences in constant dimensions. Our algorithms do not require the preferences to be narcissistic.

It is worth noting that all of our verification and hardness results apply to the stable roommates problem as well. Since finding a stable roommate matching is strictly harder than finding a stable matching, this is also optimal. Likewise, verification is equally hard for both stable roommates and stable matching, as we can simply duplicate every participant and treat the roommate matching as bipartite. Therefore, our results show that verification can be done more efficiently for the stable roommates problem when the preferences are succinct.

Finally, we address the issue of strategic behavior in these restricted models. It is often preferable for a market-clearing mechanism to incentivize truthful behavior from

the participants so that the outcome faithfully captures the optimal solution. Particularly in matching markets, this objective complements the desire for a stable matching where participants have incentives to cooperate with the outcome. Roth [74] showed that there is no strategy proof mechanism to find a stable matching in the general preferences setting. Additionally, if a mechanism outputs the man-optimal stable matching, the women can manipulate it to obtain the woman-optimal solution by truncating their preference lists [74, 34]. Even if the women are required to rank all men, they can still achieve more preferable outcomes in some instances [85, 57]. However, in the d -attribute, d -list, single-peaked, and geometric preference models, there are considerably fewer degrees of freedom for preference misrepresentation. Nevertheless, we show that there is still no strategy proof mechanism to find a stable matching for any of these models with $d \geq 2$ and non-narcissistic preferences.

Dabney and Dean [28] study an alternative succinct preference representation where there is a canonical preference list for each side and individual deviations from this list are specified separately. They provide an adaptive $O(n+k)$ time algorithm for the special one-sided case, where k is the number of deviations.

4.1 Preliminaries

A matching market (S, P) in the d -attribute model consists of n men and n women as before. A participant $s \in S$ has attributes $A_i(s)$ for $1 \leq i \leq d$ and weights $\alpha_i(s)$ for $1 \leq i \leq d$. For a man $m \in M$ and woman $w \in W$, m 's value of w is given by $\text{val}_m(w) = \langle \alpha(m), A(w) \rangle = \sum_{i=1}^d \alpha_i(m) A_i(w)$. m ranks the women in decreasing order of value. Symmetrically, w 's value of m is $\text{val}_w(m) = \sum_{i=1}^d \alpha_i(w) A_i(m)$. Note that representing a matching market in the d -attribute model requires size $O(dn)$. Unless otherwise specified, both attributes and weights can be positive or negative.

A matching market in the d -list model is a matching market where both sides

have at most d distinct preference lists. Describing a matching market in this model requires $O(dn)$ numbers.

Throughout this chapter, we use \tilde{O} to suppress polylogarithmic factors in the time complexity.

4.2 Finding Stable Matchings

4.2.1 Small Set of Attributes and Weights

We first present a stable matching algorithm for the d -attribute model when the attribute and weight values are limited to a set of constant size. In particular, we assume that the number of possible values for each attribute and weight for all participants is bounded by a constant C .

Algorithm 4.1: Small Constant Attributes and Weights

Group the women into sets S_i with a set for each of the $C' = O(C^{2d})$ types of women. ($O(C^d)$ possible attribute values and $O(C^d)$ possible weight vectors.)

Associate an empty min-heap h_i with each set S_i .

for each man m do

 Create m 's preference list of sets S_i .

 index(m) \leftarrow 1

while there is a man m who is not in any heap do

 Let S_i be the index(m) set on m 's list.

if $|h_i| < |S_i|$ **then**

h_i .insert(m)

else

if $\text{val}_{S_i}(m) > \text{val}_{S_i}(h_i.\text{min})$ **then**

h_i .delete_min()

h_i .insert(m)

 index(m) \leftarrow index(m) + 1

for $i = 1$ to C' do

$\mu \leftarrow \mu \cup$ Arbitrarily pair women in S_i with men in h_i .

return μ

Theorem 4.1. *There is an algorithm to find a stable matching in the d -attribute model with at most a constant C distinct attribute and weight values in time $O(C^{2d}n(d + \log n))$.*

Proof. Consider Algorithm 4.1. First observe that each man is indifferent between the women in a given set S_i because each woman has identical attribute values. Moreover, the women in a set S_i share the same ranking of the men, since they have identical weight vectors. Therefore, since we are looking for a stable matching, we can treat each set of women S_i as an individual entity in a many to one matching where the capacity for each S_i is the number of women it contains.

With these observations, the stability follows directly from the stability of the standard deferred acceptance algorithm for many-one stable matching. Indeed, each man proposes to the sets of women in the order of his preferences and each set of women tentatively accepts the best proposals, holding onto no more than the available capacity.

The grouping of the women requires $O(C^{2d} + dn)$ time to initialize the groups and place each woman in the appropriate group. Creating the men's preference lists requires $O(dC^{2d}n)$ time to evaluate and sort the groups of women for every man. The while loop requires $O(C^{2d}n(d + \log n))$ time since each man will propose to at most C^{2d} sets of women and each proposal requires $O(d + \log n)$ time to evaluate and update the heap. This results in an overall running time of $O(C^{2d}n(d + \log n))$. \square

As long as $d < \frac{1}{2 \log C} \log n$, the time complexity in Theorem 4.1 will be sub-quadratic. It is worth noting that the algorithm and proof actually do not rely on any restriction of the men's attribute and weight values. Thus, this result holds whenever one side's attributes and weight values come from a set of constant size.

4.2.2 One-Sided Real Attributes

In this section we consider a one-sided attribute model with real attributes and weights. In this model, women have d attributes and men have d weights, and the

preference list of a man is given by the weighted sum of the women's attributes as in the two-sided attribute model. On the other hand there is only one attribute for the men. The women's preferences are thus determined by whether they have a positive or negative weight on this attribute. For simplicity, we first assume that all women have a positive weight on the men's attribute and show a subquadratic algorithm for this case. Then we extend it to allow for negative weights.

To find a stable matching when the women have a global preference list over the men, we use a greedy approach: process the men from the most preferred to the least preferred and match each man with the highest unmatched woman in his preference list. This general technique is not specific to the attribute model but actually works for any market where one side has a single global preference list. (e.g. [28] uses a similar approach for their algorithm.) The complexity lies in repeatedly finding which of the available women is most preferred by the current top man.

This leads us to the following algorithm: for every woman w consider a point with $A(w)$ as its coordinates and organize the set of points into a data structure. Then, for the men in order of preference, query the set of points against a direction vector consisting of the man's weight and find the point with the largest distance along this direction. Remove that point and repeat.

The problem of finding a maximal point along a direction is typically considered in its dual setting, where it is called the *ray shooting problem*. In the ray shooting problem we are given n hyperplanes and must maintain a data structure to answer queries. Each query consists of a vertical ray and the data structure returns the first hyperplane hit by that ray.

The relevant results are in Lemma 4.1 which follows from several papers for different values of d . For an overview of the ray shooting problem and related range query problems, see [5].

Lemma 4.1 ([39, 30, 4, 65]). *Given an n point set in \mathbb{R}^d for $d \geq 2$, there is a data structure for ray shooting queries with preprocessing time $\tilde{O}(n)$ and query time $\tilde{O}(n^{1-1/\lfloor d/2 \rfloor})$. The structure supports deletions with amortized update time $\tilde{O}(1)$.*

For $d = 1$, queries can trivially be answered in constant time. We use this data structure to provide an algorithm when there is a global list for one side of the market.

Lemma 4.2. *For $d \geq 2$ there is an algorithm to find a stable matching in the one-sided d -attribute model with real-valued attributes and weights in time $\tilde{O}(n^{2-1/\lfloor d/2 \rfloor})$ when there is a single preference list for the other side of the market.*

Proof. For a man m , let $\dim(m)$ denote the index of the last non-zero weight, that is $\alpha_{\dim(m)+1}(m) = \dots = \alpha_d(m) = 0$. We assume $\dim(m) > 0$, as otherwise m is indifferent among all women and we can pick any woman as $\mu(m)$. We assume without loss of generality $\alpha_{\dim(m)}(m) \in \{-1, 1\}$. For each d' such that $1 \leq d' \leq d$ we build a data structure consisting of n hyperplanes in $\mathbb{R}^{d'}$. For each woman w , consider the hyperplanes

$$H_{d'}(w) = \left\{ x_{d'} = \sum_{i=1}^{d'-1} A_i(w)x_i - A_{d'}(w) \right\} \quad (4.1)$$

and for each d' preprocess the set of all hyperplanes according to Lemma 4.1. Note that $H_{d'}(w)$ is the dual of the point $(A_1(w), \dots, A_{d'}(w))$.

For a man m we can find his most preferred partner by querying the $\dim(m)$ -dimensional data structure. Let $s = \alpha_{\dim(m)}(m)$. Consider a ray $r(m) \in \mathbb{R}^{\dim(m)}$ originating at

$$\left(-\frac{\alpha_1(m)}{s}, \dots, -\frac{\alpha_{\dim(m)-1}(m)}{s}, -s \cdot \infty \right) \quad (4.2)$$

in the direction $(0, \dots, 0, s)$. If $\alpha_{\dim(m)} = 1$ we find the lowest hyperplane intersecting the ray, and if $\alpha_{\dim(m)} = -1$ we find the highest hyperplane. We claim that the first hyperplane

$r(m)$ hits corresponds to m 's most preferred woman. Let woman w be preferred over woman w' , i.e. $\text{val}_m(w) = \sum_{i=1}^{\dim(m)} A_i(w)\alpha_i(m) \geq \sum_{i=1}^{\dim(m)} A_i(w')\alpha_i(m) = \text{val}_m(w')$. Since the ray $r(m)$ is vertical in coordinate $x_{d'}$, it is sufficient to evaluate the right-hand side of the definition in equation 4.1. Indeed we have $\text{val}_m(w) \geq \text{val}_m(w')$ if and only if

$$\sum_{i=1}^{\dim(m)-1} -A_i(w)\frac{\alpha_i(m)}{s} - A_{\dim(m)}(w) \leq \sum_{i=1}^{\dim(m)-1} -A_i(w')\frac{\alpha_i(m)}{s} - A_{\dim(m)}(w') \quad (4.3)$$

when $s = 1$ and

$$\sum_{i=1}^{\dim(m)-1} -A_i(w)\frac{\alpha_i(m)}{s} - A_{\dim(m)}(w) \geq \sum_{i=1}^{\dim(m)-1} -A_i(w')\frac{\alpha_i(m)}{s} - A_{\dim(m)}(w') \quad (4.4)$$

when $s = -1$.

Note that the query ray is dual to the set of hyperplanes with normal vector $(\alpha_1(m), \dots, \alpha_d(m))$.

Now we pick the highest man m in the (global) preference list, consider the ray as above and find the first hyperplane $H_{\dim(m)}(w)$ hit by the ray. We then match the pair (m, w) , remove $H(w)$ from all data structures and repeat. Correctness follows from the correctness of the greedy approach when all women share the same preference list and the properties of the halfspaces proved above.

The algorithm preprocesses d data structures, then makes n queries and dn deletions. The time is dominated by the n ray queries each requiring time $\tilde{O}(n^{1-1/\lfloor d/2 \rfloor})$. Thus the total time complexity is bounded by $\tilde{O}(n^{2-1/\lfloor d/2 \rfloor})$, as claimed. \square

Note that for $d = 1$ there is a trivial linear time algorithm for the problem.

We use the following lemma to extend the above algorithm to account for positive and negative weights for the women. It deals with settings where the women choose one of two lists (σ_1, σ_2) as their preference lists over the men while the men's preferences

Algorithm 4.2: One-Sided Stable Matching

```

// For points in  $P \in \mathbb{R}^d$  we use the notation  $(x_1, \dots, x_d)$  to refer to its
// coordinates.
Input: matching  $\mu$ 
for  $d' = 1$  to  $d$  do
  for each woman  $w$  do
     $H(w) \leftarrow \{x_d = \sum_{i=1}^{d-1} A_i(w)x_i - A_{d'}(w)\}$ 
     $H_{d'} \leftarrow H_{d'} \cup H(w)$ 
   $H_{d'}.preprocess()$ 
for each man  $m$  in order of preference do
   $s \leftarrow \alpha_{\dim(m)}(m)$ 
   $r(m) \leftarrow (-\frac{\alpha_1(m)}{s}, \dots, -\frac{\alpha_{\dim(m)-1}(m)}{s}, \infty \cdot s) + t \cdot (0, \dots, 0, -s)$ 
   $H(w) \leftarrow \text{Query}(H_{\dim(m)}, r(m))$ 
   $\mu \leftarrow \mu \cup (m, w)$ 
  for  $d' = 1$  to  $d$  do
     $H_{d'} \leftarrow H_{d'} - H_{d'}(w)$ 
return  $\mu$ 

```

can be arbitrary.

Lemma 4.3. *Suppose there are k women who use σ_1 . If the top k men in σ_1 are in the bottom k places in σ_2 , then the women using σ_1 will only match with those men and the $n - k$ women using σ_2 will only match with the other $n - k$ men in the woman-optimal stable matching.*

Proof. Consider the operation of the woman-proposing deferred acceptance algorithm for finding the woman-optimal stable matching. Suppose the lemma is false so that at some point a woman using σ_1 proposed to one of the last $n - k$ men in σ_1 . Let w be the first such woman. w must have been rejected by all of the top k , so at least one of those men received a proposal from a woman, w' , using σ_2 . However, since the top k men in σ_1 are the bottom k men in σ_2 , w' must have been rejected by all of the top $n - k$ men in σ_2 . But there are only $n - k$ women using σ_2 , so one of the top $n - k$ men in σ_2 must have already received a proposal from a woman using σ_1 . This is a contradiction because

Table 4.1. Two-list preferences where no participant receives their top choice in the stable matching

σ_1	σ_2	π_1	π_2	Man	List	Woman	List
m_1	m_3	w_1	w_3	m_1	π_1	w_1	σ_2
m_2	m_5	w_2	w_5	m_2	π_1	w_2	σ_2
m_3	m_1	w_3	w_1	m_3	π_2	w_3	σ_1
m_4	m_4	w_4	w_4	m_4	π_1	w_4	σ_2
m_5	m_2	w_5	w_2	m_5	π_2	w_5	σ_1

w was the first woman using σ_1 to propose to one of the bottom $n - k$ men in σ_1 (which are the top $n - k$ men in σ_2). \square

We can now prove the following theorem where negative values are allowed for the women's weights.

Theorem 4.2. *For $d \geq 2$ there is an algorithm to find a stable matching in the one-sided d -attribute model with real-valued attributes and weights in time $\tilde{O}(n^{2-1/\lfloor d/2 \rfloor})$.*

Proof. Suppose there are k women who have a positive weight on the men's attribute. Since the remaining $n - k$ women's preference list is the reverse, we can use Lemma 4.3 to split the problem into two subproblems. Namely, in the woman-optimal stable matching the k women with a positive weight will match with the top k men, and the $n - k$ women with a negative weight will match with the bottom $n - k$ men. Now the women in each of these subproblems all have the same list. Therefore we can use Lemma 4.2 to solve each subproblem. Splitting the problem into subproblems can be done in time $O(n)$ so the running time follows immediately from Lemma 4.2. \square

As a remark, this "greedy" approach where we select a man, find his most preferred available woman, and permanently match him to her will not work in general. Table 4.1 describes a simple 2-list example where the unique stable matching

Table 4.2. Two-list preferences where a greedy approach will not work

σ_1	σ_2	π_1	π_2	Man	List	Woman	List
m_1	m_2	w_1	w_3	m_1	π_2	w_1	σ_1
m_2	m_1	w_2	w_2	m_2	π_1	w_2	σ_2
m_3	m_3	w_3	w_1	m_3	π_1	w_3	σ_2

is $\{(m_1, w_2), (m_2, w_3), (m_3, w_5), (m_4, w_4), (m_5, w_1)\}$. In this instance, no participant is matched with their top choice. Therefore, the above approach cannot work for this instance. This illustrates to some extent why the general case seems more difficult than the one-sided case.

An alternative model of a greedy approach that is based on work by Davis and Impagliazzo in [29] also will not work. In this model, an algorithm can view each of the lists and the preferences of the women. It can then (adaptively) choose an order in which to process the men. When processing a man, he must be assigned a partner (not necessarily his favorite available woman) once and for all, based only on his choice of preference list and the preferences of the previously processed men. This model is similar to online stable matching [55] except that it allows the algorithm to choose the processing order of the men. Using the preferences in Table 4.2 and minor modifications to them, we can show that no greedy algorithm of this type can successfully produce a stable matching. Indeed, the unique stable matching of the preference scheme below is $\mu = \{(m_1, w_3), (m_2, w_1), (m_3, w_2)\}$. However, changing the preference list for whichever of m_1 or m_2 is processed later will form a blocking pair with the stable partner of the other. If m_1 uses π_1 , (m_1, w_1) blocks μ and if m_2 uses π_2 , (m_2, w_3) blocks μ . Therefore, no algorithm can succeed in assigning stable partners to these men without first knowing the preference list choice of all three.

4.2.3 Strategic Behavior

As mentioned earlier, strategic behavior in the general preference setting allows for participants to truncate or rearrange their lists. However, in the d -attribute and d -list models, we assume that the attributes or lists are fixed, so that the only manipulation the participants are allowed is to misrepresent their weight vectors or which list they choose. Despite this limitation, there is still no strategy proof mechanism for finding a stable matching when $d \geq 2$.

Theorem 4.3. *For $d \geq 2$ there is no strategy proof algorithm to find a stable matching in the d -list model.*

Proof. Table 4.3 describes true preferences that can be manipulated by the women. Observe that there are two stable matchings: $\{(m_1, w_1), (m_2, w_2), (m_3, w_3), (m_4, w_4)\}$, the man-optimal matching, and $\{(m_1, w_2), (m_2, w_3), (m_3, w_1), (m_4, w_4)\}$, the woman-optimal matching. However, if w_2 used list σ_2 instead of σ_1 , then there is a unique stable matching which is $\{(m_1, w_2), (m_2, w_3), (m_3, w_1), (m_4, w_4)\}$, the woman-optimal stable matching from the original preferences. Therefore, any mechanism that does not always output the woman optimal stable matching can be manipulated by the women to their advantage. By symmetry, any mechanism that does not always output the man-optimal matching could be manipulated by the men. Thus there is no strategy-proof mechanism for the d -list setting with $d \geq 2$. \square

Since the d -list model is a special case of the d -attribute model, we immediately have the following result from Theorem 4.3.

Corollary 4.1. *For $d \geq 2$ there is no strategy proof algorithm to find a stable matching in the d -attribute model.*

Table 4.3. Two-list preferences that can be manipulated

σ_1	σ_2	π_1	π_2	Man	List	Woman	List
m_1	m_3	w_1	w_3	m_1	π_1	w_1	σ_2
m_2	m_1	w_2	w_1	m_2	π_1	w_2	σ_1
m_3	m_4	w_3	w_2	m_3	π_2	w_3	σ_1
m_4	m_2	w_4	w_4	m_4	π_2	w_4	σ_1

Of course in the 1-list setting there is a trivial unique stable matching. Moreover, in the one-sided d -attribute model our algorithm is strategy proof since the women are receiving the woman-optimal matching and each man receives his best available woman, so misrepresentation would only give him a worse partner.

4.3 Verification

We now turn to the problem of verifying whether a given matching is stable. While this is as hard as finding a stable matching in the general setting, the verification algorithms we present here are more efficient than our algorithms for finding stable matchings in the attribute model.

4.3.1 Real Attributes and Weights

In this section we adapt the geometric approach for finding a stable matching in the one-sided d -attribute model to the problem of verifying a stable matching in the (two-sided) d -attribute model. We express the verification problem as a *simplex range searching problem* in \mathbb{R}^{2d} , which is the dual of the ray shooting problem. In simplex range searching we are given n points and answer queries that ask for the number of points inside a simplex. In our case we only need degenerate simplices consisting of the intersection of two halfspaces. Simplex range searching queries can be done in sublinear time for constant d .

Lemma 4.4 ([64]). *Given a set of n points in \mathbb{R}^d , one can process it for simplex range searching in time $O(n \log n)$, and then answer queries in time $\tilde{O}(n^{1-\frac{1}{d}})$.*

For $1 \leq d' \leq d$ we use the notation $(x_1, \dots, x_d, y_1, \dots, y_{d'-1}, z)$ for points in $\mathbb{R}^{d+d'}$. We again let $\dim(w)$ be the index of w 's last non-zero weight, assume without loss of generality $\alpha_{\dim(w)} \in \{-1, 1\}$, and let $\text{sgn}(w) = \text{sgn}(\alpha_{\dim(w)})$. We partition the set of women into $2d$ sets $W_{d',s}$ for $1 \leq d' \leq d$ and $s \in \{-1, 1\}$ based on $\dim(w)$ and $\text{sgn}(w)$. Note that if $\dim(w) = 0$, then w is indifferent among all men and can therefore not be part of a blocking pair. We can ignore such women.

For a woman w , consider the point

$$P(w) = (A_1(w), \dots, A_d(w), \alpha_1(w), \dots, \alpha_{\dim(w)-1}(w), \text{val}_w(m)) \quad (4.5)$$

where $m = \mu(w)$ is the partner of w in the input matching μ . For a set $W_{d',s}$ we let $P_{d',s}$ be the set of points $P(w)$ for $w \in W_{d',s}$. The basic idea is to construct a simplex for every man and query it against all sets $P_{d',s}$.

Given d',s , and a man m , let $H_1(m)$ be the halfspace $\{\sum_{i=1}^d \alpha_i(m)x_i > \text{val}_m(w)\}$ where $w = \mu(m)$. For $w' \in W_{d',s}$ we have $P(w') \in H_1(m)$ if and only if m strictly prefers w' to w . Further let $H_2(m)$ be the halfspace $\{\sum_{i=1}^{d'-1} A_i(m)y_i + A_{d'}(m)s > z\}$. For $w' \in W_{d',s}$ we have $P(w') \in H_2(m)$ if and only if w' strictly prefers m to $\mu(w')$. Hence (m, w') is a blocking pair if and only if $P(w') \in H_1(m) \cap H_2(m)$.

Using Lemma 4.4 we immediately have an algorithm to verify a stable matching.

Theorem 4.4. *There is an algorithm to verify a stable matching in the d -attribute model with real-valued attributes and weights in time $\tilde{O}(n^{2-1/2d})$*

Proof. Partition the set of women into sets $W_{d',s}$ for $1 \leq d' \leq d$ and $s \in \{-1, 1\}$ and for $w \in W_{d',s}$ construct $P(w) \in \mathbb{R}^{d+d'}$ as above. Then preprocess the sets according to

Lemma 4.4. For each man m query $H_1(m) \cap H_2(m)$ against the points in all sets. By the definitions of $H_1(m)$ and $H_2(m)$, there is a blocking pair if and only if for some man m there is a point $P(w) \in H_1(m) \cap H_2(m)$ in one of the sets $P_{d',s}$.

The time to preprocess is $O(n \log n)$. There are $2dn$ queries of time $\tilde{O}(n^{1-1/2d})$. Hence the whole process requires time $\tilde{O}(n^{2-1/2d})$ as claimed. \square

Algorithm 4.3: Verify Stable Matching with Reals

// For points in $P \in \mathbb{R}^{d+d'}$ we use the notation $(x_1, \dots, x_d, y_1, \dots, y_{d'-1}, z)$ to refer to its coordinates.

Input: matching μ

for each woman w do

$m \leftarrow \mu(w)$

$P(w) \leftarrow (A_1(w), \dots, A_d(w), \alpha_1(w), \dots, \alpha_d(w), \text{val}_w(m))$

$P_{\dim(w), \text{sgn}(w)} \leftarrow W_{\dim(w), \text{sgn}(w)} \cup P(w)$

for $d' = 1$ to d and $s \in \{-1, 1\}$ do

$P_{d',s}.\text{preprocess}()$

for each man m do

$w \leftarrow \mu(m)$

$H_1(m) \leftarrow \{ \sum_{i=1}^d \alpha_i(m)x_i > \text{val}_m(w) \}$

$H_2(m) \leftarrow \{ \sum_{i=1}^{d'-1} A_i(m)y_i + A_{d'}(m) \cdot s > z \}$

if Query($P_{d',s}, H_1(m) \cap H_2(m)$) > 0 then

return μ is not stable

return μ is stable

4.3.2 Lists

When there are d preference orders for each side, and each participant uses one of the d lists, we provide a more efficient algorithm. Here, assume μ is the given matching between M and W . Let $\{\pi_i\}_{i=1}^d$ be the set of d permutations on the women and $\{\sigma_j\}_{j=1}^d$ be the set of d permutations on the men. Define $\text{rank}(w, i)$ to be the position of w in permutation π_i . This can be determined in constant time after $O(dn)$ preprocessing of the permutations. Let $\text{head}(\pi_i, j)$ be the first woman in π_i who uses permutation σ_j and

$\text{next}(w, i)$ be the next highest ranked woman after w in permutation π_i who uses the same permutation as w or \perp if no such woman exists. These can also be determined in constant time after $O(dn)$ preprocessing by splitting the lists into sublists, with one sublist for the women using each permutation of men. The functions rank , head , and next are defined analogously for the men.

Algorithm 4.4: Verify d -List Stable Matching

```

for  $i = 1$  to  $d$  do
  for  $j = 1$  to  $d$  do
     $w \leftarrow \text{head}(\pi_i, j)$ .
     $m \leftarrow \text{head}(\sigma_j, i)$ .
    while  $m \neq \perp$  and  $w \neq \perp$  do
      if  $\text{rank}(w, i) > \text{rank}(\mu(m), i)$  then
         $m \leftarrow \text{next}(m, j)$ .
      else
        if  $\text{rank}(m, j) > \text{rank}(\mu(w), j)$  then
           $w \leftarrow \text{next}(w, i)$ .
        else
          return  $(m, w)$  is a blocking pair.
  return  $\mu$  is stable.

```

Theorem 4.5. *There is an algorithm to verify a stable matching in the d -list model in $O(dn)$ time.*

Proof. We claim that algorithm 4.4 satisfies the theorem. Indeed, if the algorithm returns a pair (m, w) where m uses π_i and w uses σ_j , then (m, w) is a blocking pair because w appears earlier in π_i than $\mu(m)$ and m appears earlier in σ_j than $\mu(w)$.

On the other hand, suppose the algorithm returns that μ is stable but there is a blocking pair, (m, w) , where m uses π_i and w uses σ_j . The algorithm considers permutations π_i and σ_j since it does not terminate early. Clearly if the algorithm evaluates m and w simultaneously when considering permutations π_i and σ_j , it will detect that

(m, w) is a blocking pair. Therefore, the algorithm either moves from m to $\text{next}(m, j)$ before considering w or it moves from w to $\text{next}(w, i)$ before considering m . In the former case, $\text{rank}(\mu(m), i) < \text{rank}(w', i)$ for some w' that comes before w in π_i . Therefore m prefers $\mu(m)$ to w . Similarly, in the latter case, $\text{rank}(\mu(w), j) < \text{rank}(m', i)$ for some m' that comes before m in σ_j so w prefers $\mu(w)$ to m . Thus (m, w) is not a blocking pair and we have a contradiction.

The **for** and **while** loops proceed through all men and women once for each of the d lists in which they appear. Since at each step we are either proceeding to the next man or the next woman unless we find a blocking pair, the algorithm requires time $O(dn)$. This is optimal since the input size is dn . \square

4.3.3 Boolean Attributes and Weights

In this section we consider the problem of verifying a stable matching when the d attributes and weights are restricted to boolean values and $d = c \log n$. The algorithm closely follows an algorithm for the maximum inner product problem by Alman and Williams [6]. The idea is to express the existence of a blocking pair as a probabilistic polynomial with a bounded number of monomials and use fast rectangular matrix multiplication to evaluate it. A probabilistic polynomial for a function f is a polynomial p such that for every input x

$$\Pr[f(x) \neq p(x)] \leq \frac{1}{3} \quad (4.6)$$

We use the following tools in our algorithm. THR_d is the threshold function that outputs 1 if at least d of its inputs are 1.

Lemma 4.5 ([6]). *There is a probabilistic polynomial for THR_d on n variables and error ε with degree $O(\sqrt{n \log(1/\varepsilon)})$.*

Lemma 4.6 ([73, 83]). *There is a probabilistic polynomial for the disjunction of n variables and error ε with degree $O(\log(1/\varepsilon))$*

Lemma 4.7 ([87]). *Given a polynomial $P(x_1, \dots, x_m, y_1, \dots, y_m)$ with at most $n^{0.17}$ monomials and two sets $X, Y \subseteq \{0, 1\}^m$ with $|X| = |Y| = n$, we can evaluate P on all pairs $(x, y) \in X \times Y$ in time $\tilde{O}(n^2 + m \cdot n^{1.17})$.*

We construct a probabilistic polynomial that outputs 1 if there is a blocking pair. To minimize the degree of the polynomial, we pick a parameter s and divide the men and women into sets of size at most s . The polynomial takes the description of s men m_1, \dots, m_s and s women w_1, \dots, w_s along with their respective partners as input, and outputs 1 if and only if there is a blocking pair (m_i, w_j) among the s^2 pairs of nodes with high probability.

Lemma 4.8. *Let u be a large constant and $s = n^{1/uc \log^2 c}$. There is a probabilistic polynomial with the following inputs:*

- *The attributes and weights of s men: $A(m_1), \dots, A(m_s), \alpha(m_1), \dots, \alpha(m_s)$*
- *The attributes of the s women matched with these men: $A(\mu(m_1)), \dots, A(\mu(m_s))$*
- *The attributes and weights of s women: $A(w_1), \dots, A(w_s), \alpha(w_1), \dots, \alpha(w_s)$*
- *The attributes of the s men matched with these women: $A(\mu(w_1)), \dots, A(\mu(w_s))$*

The output of the polynomial is 1 if and only if there is a blocking pair with respect to the matching μ among the s^2 pairs in the input. The number of monomials is at most $n^{0.17}$ and the polynomial can be constructed efficiently.

Proof. (m_i, w_j) is a blocking pair if and only if $\langle \alpha(m_i), A(\mu(m_i)) \rangle < \langle \alpha(m_i), A(w_j) \rangle$ and

$\langle \alpha(w_j), A(\mu(w_j)) \rangle < \langle \alpha(w_j), A(m_i) \rangle$. Rewriting

$$F(x, y, a, b) := \langle x, y \rangle < \langle a, b \rangle = \text{THR}_{d+1}(\neg(x_1 \wedge y_1), \dots, \neg(x_d \wedge y_d), a_1 \wedge b_1, \dots, a_d \wedge b_d) \quad (4.7)$$

we have a blocking pair if and only if

$$\bigvee_{\substack{i \in [1, s] \\ j \in [1, s]}} (F(\alpha(m_i), A(\mu(m_i)), \alpha(m_i), A(w_j)) \wedge F(\alpha(w_j), A(\mu(w_j)), \alpha(w_j), A(m_i))) \quad (4.8)$$

Note that we can easily adapt this algorithm to finding strongly blocking pairs by defining $F(x, y, a, b)$ as $\langle x, y \rangle \leq \langle a, b \rangle$.

Using Lemma 4.5 with $\varepsilon = \frac{1}{s^3}$ and Lemma 4.6 with $\varepsilon = 1/4$ we get a probabilistic polynomial of degree $a\sqrt{d \log s}$ for some constant a and error $1/4 + 1/s < 1/3$. Furthermore, since we are only interested in boolean inputs we can assume the polynomial to be multilinear. For large enough u we have $2d > a\sqrt{d \log(s)}$ (i.e. the degree is at most half of the number of variables) and the number of monomials is then bounded by $O\left(\left(s^2 \binom{4d}{a\sqrt{d \log(s)}}\right)^2\right)$.

Simplifying the binomial coefficient we have

$$\binom{4d}{a\sqrt{d \log s}} = \binom{4c \log n}{a\sqrt{(\log^2 n)/u \log^2 c}} = \binom{4c \log n}{a \log n / \sqrt{u} \log c}$$

Setting $\delta = a/(\sqrt{u} \log(c))$ we can upper bound this using Stirling's inequality by

$$\binom{4c \log n}{\delta \log n} \leq \left(\frac{(4c \log n) \cdot e}{\delta}\right)^{\delta \log n} = n^{\delta \log(4ce/\delta)}$$

By choosing u to be a large enough constant, we can make δ and the exponent arbitrarily small. The factor of s^2 only contributes a trivial constant to the exponent.

Therefore we can bound the number of monomials by $n^{0.17}$. \square

Theorem 4.6. *In the d -attribute model with n men and women, and $d = c \log n$ boolean attributes and weights, there is a randomized algorithm to decide if a given matching is stable in time $\tilde{O}(n^{2-1/O(c \log^2(c))})$ with error probability at most $1/3$.*

Proof. We again choose $s = n^{1/uc \log^2 c}$ and construct the probabilistic polynomial as in Lemma 4.8. We then divide the men and women into $\lceil \frac{n}{s} \rceil$ groups of size at most s .

For a group of men m_1, \dots, m_s we let the corresponding input vector be

$$A(m_1), \dots, A(m_s), \alpha(m_1), \dots, \alpha(m_s), A(\mu(m_1)), \dots, A(\mu(m_s))$$

We set X as the set of all input vectors for the $\lceil \frac{n}{s} \rceil$ groups. We define the set Y symmetrically for the input vectors corresponding to the $\lceil \frac{n}{s} \rceil$ groups of women.

Using Lemma 4.7 we evaluate the polynomial on all pairs $x \in X, y \in Y$ in time

$$\tilde{O} \left(\left(\frac{n}{s} \right)^2 + O(sd) \left(\frac{n}{s} \right)^{1.17} \right) = \tilde{O} \left(\left(\frac{n}{s} \right)^2 \right) = \tilde{O}(n^{2-1/O(c \log^2(c))}) \quad (4.9)$$

The probability that the output is wrong for any fixed input pair is at most $1/3$. We repeat this process $O(\log n)$ times and take the threshold output for every pair of inputs, such that the error probability is at most $O\left(\frac{1}{n^2}\right)$ for any fixed pair of inputs. Using a union bound we can make the probability of error at most $1/3$ on any input. \square

4.4 Conditional Hardness

4.4.1 Background

The Strong Exponential Time Hypothesis has proved useful in arguing conditional hardness for a large number of problems. We show SETH-hardness for both verifying

and finding a stable matching in the d -attribute model, even if the weights and attributes are boolean. The main step of the proof is a reduction from the maximum inner product problem to the stable matching problem. The maximum inner product problem is known to be SETH-hard. We give the fine-grained reduction from CNFSAT to the vector orthogonality problem and from the vector orthogonality problem to the maximum inner product problem for the sake of completeness.

Definition 4.1 ([45, 47]). *The Strong Exponential Time Hypothesis (SETH) stipulates that for each $\varepsilon > 0$ there is a k such that k -SAT requires time $\Omega(2^{(1-\varepsilon)n})$.*

Definition 4.2. *For any d , the vector orthogonality problem is to decide if two input sets $U, V \subseteq \mathbb{R}^d$ with $|U| = |V| = n$ have a pair $u \in U, v \in V$ such that $\langle u, v \rangle = 0$.*

The boolean vector orthogonality problem is the variant where $U, V \subseteq \{0, 1\}^d$.

Definition 4.3. *For any d and input l , the maximum inner product problem is to decide if two input sets $U, V \subseteq \mathbb{R}^d$ with $|U| = |V| = n$ have a pair $u \in U, v \in V$ such that $\langle u, v \rangle \geq l$.*

The boolean maximum inner product problem is the variant where $U, V \subseteq \{0, 1\}^d$.

Lemma 4.9 ([47, 86, 6]). *Assuming SETH, for any $\varepsilon > 0$, there is a c such that solving the boolean maximum inner product problem on $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. The proof is a series of reductions from k -SAT to boolean inner product. By the Sparsification Lemma [47] we can reduce k -SAT to a subexponential number of k -SAT instances with at most $d = c_k n$ clauses, where c_k does not depend on n . Hence, assuming SETH, for any $\varepsilon > 0$, there is a c such that CNFSAT with cn clauses requires time $\Omega(2^{(1-\varepsilon)n})$.

We reduce CNFSAT to the boolean vector orthogonality problem using a technique called *Split and List*. Divide the variable set into two sets S, T of size $\frac{n}{2}$ and for each set

consider all $N = 2^{n/2}$ assignments to the variables. For every assignment we construct a d -dimensional vector where the i th position is 1 if and only if the assignment does not satisfy the i th clause of the CNF formula. Let U be the set of vectors corresponding to the assignments to S and let V be the set of vectors corresponding to T . A pair $u \in U$, $v \in V$ is orthogonal if and only if the corresponding assignment satisfies all clauses. An algorithm for boolean vector orthogonality in dimension $d = cn = 2c \log N$ and time $O(N^{2-\varepsilon}) = O(2^{(1-\varepsilon/2)n})$ would contradict SETH. Hence assuming SETH, for every $\varepsilon > 0$ there is a c such that the boolean vector orthogonality problem with $d = c \log N$ requires time $\Omega(N^{2-\varepsilon})$.

Finally, we reduce the boolean vector orthogonality problem to the boolean maximum inner product problem by partitioning the set U into sets U_i for $0 \leq i \leq d$ where U_i contains all vectors with Hamming weight i . Observe that a vector $v \in V$ is orthogonal to a vector $u \in U_i$ if and only if $\langle u, \neg v \rangle = i$, where $\neg v$ is the element-wise complement of v . Thus U and V have an orthogonal pair, if and only if there is an i such that U_i and $\neg V = \{\neg v \mid v \in V\}$ have a pair with inner product at least i . Therefore, for any $\varepsilon > 0$ there is a c such that the maximum inner product problem on $d = c \log N$ dimensions requires time $\Omega(N^{2-\varepsilon})$ assuming SETH. \square

4.4.2 Finding Stable Matchings

In this subsection we give a fine-grained reduction from the maximum inner product problem to the problem of finding a stable matching in the boolean d -attribute model. This shows that the stable matching problem in the d -attribute model is SETH-hard, even if we restrict the attributes and weights to booleans.

Theorem 4.7. *Assuming SETH, for any $\varepsilon > 0$, there is a c such that finding a stable matching in the boolean d -attribute model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. The proof is a reduction from maximum inner product to finding a stable matching. Given an instance of the maximum inner product problem with sets $U, V \subseteq \{0, 1\}^d$ where $|U| = |V| = n$ and threshold l , we construct a matching market with n men and n women. For every $u \in U$ we have a man m_u with $A(m_u) = u$ and $\alpha(m_u) = u$. Similarly, for vectors $v \in V$ we have women w_v with $A(w_v) = v$ and $\alpha(w_v) = v$. This matching market is symmetric in the sense that for m_u and w_v , $\text{val}_{m_u}(w_v) = \text{val}_{w_v}(m_u) = \langle u, v \rangle$.

We claim that any stable matching contains a pair (m_u, w_v) such that the inner product $\langle u, v \rangle$ is maximized. Indeed, suppose there are vectors $u \in U, v \in V$ with $\langle u, v \rangle \geq l$ but there exists a stable matching μ with $\langle u', v' \rangle < l$ for all pairs $(m_{u'}, w_{v'}) \in \mu$. Then (m_u, w_v) is clearly a blocking pair for μ which is a contradiction. \square

4.4.3 Verifying Stable Matchings

In this section we give a reduction from the maximum inner product problem to the problem of verifying a stable matching, showing that this problem is also SETH-hard.

Theorem 4.8. *Assuming SETH, for any $\varepsilon > 0$, there is a c such that verifying a stable matching in the boolean d -attribute model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. We give a reduction from maximum inner product with sets $U, V \subseteq \{0, 1\}^d$ where $|U| = |V| = n$ and threshold l . We construct a matching market with $2n$ men and women in the d' -attribute model with $d' = d + 2(l - 1)$. Since $d' < 3d$ the theorem then follows immediately from the SETH-hardness of maximum inner product.

For $u \in U$, let m_u be a man in the matching market with attributes and weights $A(m_u) = \alpha(m_u) = u \circ 1^{l-1} \circ 0^{l-1}$ where we use \circ for concatenation. Similarly, for $v \in V$ we have a woman w_v with $A(w_v) = \alpha(w_v) = v \circ 0^{l-1} \circ 1^{l-1}$. We further introduce *dummy women* w'_u for $u \in U$ with $A(w'_u) = \alpha(w'_u) = 0^d \circ 1^{l-1} \circ 0^{l-1}$ and *dummy men* m'_v for $v \in V$ with $A(m'_v) = \alpha(m'_v) = 0^d \circ 0^{l-1} \circ 1^{l-1}$.

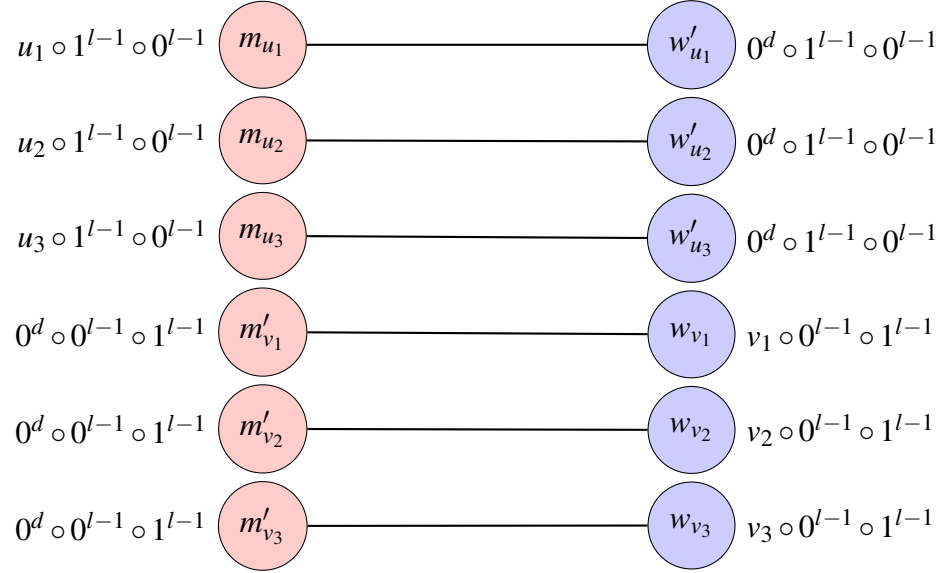


Figure 4.1. A representation of the reduction from maximum inner product to verifying a stable matching

We claim that the matching consisting of pairs (m_u, w'_u) for all $u \in U$ and (m'_v, w_v) for all $v \in V$ is stable if and only if there is no pair $u \in U, v \in V$ with $\langle u, v \rangle \geq l$. For $u, u' \in U$ we have $\text{val}_{m_u}(w'_{u'}) = \text{val}_{w'_{u'}}(m_u) = l - 1$, and for $v, v' \in V$ we have $\text{val}_{w_v}(m'_{v'}) = \text{val}_{m'_{v'}}(w_v) = l - 1$. In particular, any pair in μ has (symmetric) value $l - 1$. Hence there is a blocking pair with respect to μ if and only if there is a pair with value at least l . For $u \neq u'$ and $v \neq v'$ the pairs $(m_u, w'_{u'})$ and $(w_v, m'_{v'})$ can never be blocking pairs as their value is $l - 1$. Furthermore for any pair of dummy nodes w'_u and m'_v we have $\text{val}_{m'_v}(w'_u) = \text{val}_{w'_u}(m'_v) = 0$, thus no such pair can be a blocking pair either. This leaves pairs of real nodes as the only candidates for blocking pairs. For non-dummy nodes m_u and w_v we have $\text{val}_{m_u}(w_v) = \text{val}_{w_v}(m_u) = \langle u, v \rangle$ so (m_u, w_v) is a blocking pair if and only if $\langle u, v \rangle \geq l$. \square

4.4.4 Checking a Stable Pair

In this section we give a reduction from the maximum inner product problem to the problem of checking whether a given pair is part of any or all stable matchings, showing that these questions are SETH-hard when $d = c \log n$ for some constant c . For general preferences, both questions can be solved in time $O(n^2)$ [50, 37] and are known to require quadratic time [36].

Theorem 4.9. *Assuming SETH, for any $\varepsilon > 0$, there is a c such that determining whether a given pair is part of any or all stable matchings in the boolean d -attribute model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. We again give a reduction from maximum inner product with sets $U, V \subseteq \{0, 1\}^d$ where $|U| = |V| = n$ and threshold l . We construct a matching market with $2n$ men and women in the d' -attribute model with $d' = 7d + 7(l - 1) + 18$. Since $d' < 15d$ the theorem then follows immediately from the SETH-hardness of maximum inner product.

For simplicity, we will first describe the preference scheme, then provide weight and attribute vectors that result in those preferences. For $u \in U$, let m_u be a man in the matching market and for $v \in V$ we have a woman w_v . We also have $n - 1$ dummy men $m_i : i = 1 \dots n - 1$ and $n - 1$ dummy women $w_j : j = 1 \dots n - 1$. Finally, we have a special man m^* and special woman w^* . This special pair is the one we will test for stability. Let the preferences be

$$\begin{aligned}
m_u : \{w_v : \langle u, v \rangle \geq l\} \succ \{w_j\}_{j=1}^{n-1} \succ w^* \succ \{w_v : \langle u, v \rangle < l\} & \quad \forall u \in U \\
m_i : \{w_v\} \succ \{w_j\}_{j=1}^{n-1} \succ w^* & \quad \forall i \in \{1 \dots n-1\} \\
m^* : w^* \succ \{w_v\} \succ \{w_j\}_{j=1}^{n-1} & \\
w_v : \{m_u : \langle u, v \rangle \geq l\} \succ \{m_i\}_{i=1}^{n-1} \succ m^* \succ \{m_u : \langle u, v \rangle < l\} & \quad \forall v \in V \\
w_j : \{m_u\} \succ \{m_i\}_{i=1}^{n-1} \succ m^* & \quad \forall j \in \{1 \dots n-1\} \\
w^* : \{m_i\}_{i=1}^{n-1} \succ \{m_u\} \succ m^* &
\end{aligned}$$

so that, for example, man m_u corresponding to $u \in U$ will most prefer women w_v for some $v \in V$ with $\langle u, v \rangle \geq l$ (in decreasing order of $\langle u, v \rangle$), then all of the dummy women (equally), then the special woman w^* , and finally the remaining women w_v (in decreasing order of $\langle u, v \rangle$).

First suppose for some $\hat{u} \in U$ and $\hat{v} \in V$ we have $\langle \hat{u}, \hat{v} \rangle \geq l$ and let this be the pair with largest inner product. Now consider the deferred acceptance algorithm for finding the woman-optimal stable matching. First, $w_{\hat{v}}$ will propose to $m_{\hat{u}}$ and will be accepted. The dummy women will propose to the remaining men corresponding to U . Then any other woman w_v will be accepted by either a dummy man or a man m_u , causing the dummy woman matched with him to move to a dummy man. In any case, all men besides m^* are matched to a woman they prefer over w^* , so when she proposes to them, they will reject her. Thus w^* will match with m^* . Since w^* receives her least preferred choice in the woman optimal stable matching, (m^*, w^*) is a pair in every stable matching.

Now suppose $\langle u, v \rangle < l$ for every $u \in U, v \in V$. Consider the deferred acceptance algorithm for finding the man-optimal stable matching. First, the dummy men will propose to the women corresponding to V and will be accepted. Then every man m_u will

propose to the dummy women, but only $n - 1$ of them can be accepted. The remaining one will propose to w^* . When m^* proposes to w^* , she rejects him, causing him to eventually be accepted by the available woman w_v . Thus m^* will not match with w^* in any stable matching since she is his most preferred choice but he is not matched with her in the man-optimal stable matching, so (m^*, w^*) is not a pair in any stable matching. Figure 4.2 demonstrates each of these cases.


Since the stable pair questions for whether (m^*, w^*) are a stable pair in any or all stable matchings are equivalent with these preferences, this reduction works for both.

Finally, we claim the following vectors realize the preferences above for the attribute model. We leave it to the reader to verify this. As in our other hardness reductions, the weight and attribute vectors are identical for each participant.

$$\begin{aligned}
m_u &: u^7 \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^6 \circ 0^6 \circ 0^6 \\
m_i &: 0^{7d} \circ 1^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 1^6 \circ 0^6 \\
m^* &: 0^{7d} \circ 1^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 0^6 \circ 1^6 \\
w_v &: v^7 \circ 0^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 1^6 \circ 1 \cdot 0^5 \\
w_j &: 0^{7d} \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^6 \circ 1^5 \circ 0 \circ 0^6 \\
w^* &: 0^{7d} \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^3 \circ 0^3 \circ 1^4 \circ 0^2 \circ 1^2 \circ 0^4
\end{aligned}$$

□

This reduction also has consequences on the existence of nondeterministic algorithms for the stable pair problem assuming the *Nondeterministic Strong Exponential Time Hypothesis* (NSETH).



(a) A representative stable matching when there is a pair (u, v) with $\langle u, v \rangle \geq l$ (b) A representative stable matching when $\langle u, v \rangle < l$ for every pair (u, v)

Figure 4.2. A representation of the reduction from maximum inner product to checking a stable pair

Definition 4.4 ([22]). *The Nondeterministic Strong Exponential Time Hypothesis stipulates that for each $\varepsilon > 0$ there is a k such that k -SAT requires co-nondeterministic time $\Omega(2^{(1-\varepsilon)n})$.*

In other words, the Nondeterministic Strong Exponential Time Hypothesis stipulates that for CNFSAT there is no proof of unsatisfiability that can be checked deterministically in time $\Omega(2^{(1-\varepsilon)n})$.

Assuming NSETH, any problem that is SETH-hard at time $T(n)$ under deterministic reductions either require $T(n)$ time nondeterministically or co-nondeterministically, i.e. either there is no proof that an instance is true or there is no proof that an instance is false that can be checked in time faster than $T(n)$. Note that all reductions in this chapter are deterministic. In particular, the maximum inner product problem does not have a $O(N^{2-\varepsilon})$ co-nondeterministic time algorithm for any $\varepsilon > 0$ assuming NSETH, since it

has a simple linear time nondeterministic algorithm.

Since the reduction of Theorem 4.9 is a simple reduction that maps a true instance of maximum inner product to a true instance of the stable pair problem, we can conclude that the stable pair problem is also hard co-nondeterministically.

Corollary 4.2. *Assuming NSETH, for any $\varepsilon > 0$, there is a c such that determining whether a given pair is part of any or all stable matchings in the boolean d -attribute model with $d = c \log n$ dimensions requires co-nondeterministic time $\Omega(n^{2-\varepsilon})$.*

We also have a reduction so that the given pair is stable in any or all stable matchings if and only if there is not a pair of vectors with large inner product. This shows that the stable pair problem is also hard nondeterministically.

Theorem 4.10. *Assuming NSETH, for any $\varepsilon > 0$, there is a c such that determining whether a given pair is part of any or all stable matchings in the boolean d -attribute model with $d = c \log n$ dimensions requires nondeterministic time $\Omega(n^{2-\varepsilon})$.*

Proof. This reduction uses the same setup as the one in Theorem 4.9 except that we now have n dummy men and women instead of $n - 1$ and we slightly change the preferences as follows:

$$\begin{aligned}
m_u : \{w_v : \langle u, v \rangle \geq l\} \succ \{w_j\}_{j=1}^n \succ w^* \succ \{w_v : \langle u, v \rangle < l\} & \quad \forall u \in U \\
m_i : \{w_v\} \succ w^* \succ \{w_j\}_{j=1}^n & \quad \forall i \in \{1 \dots n\} \\
m^* : w^* \succ \{w_v\} \succ \{w_j\}_{j=1}^n & \\
w_v : \{m_u : \langle u, v \rangle \geq l\} \succ \{m_i\}_{i=1}^n \succ m^* \succ \{m_u : \langle u, v \rangle < l\} & \quad \forall v \in V \\
w_j : \{m_u\} \succ \{m_i\}_{i=1}^n \succ m^* & \quad \forall j \in \{1 \dots n\} \\
w^* : \{m_i\}_{i=1}^n \succ \{m_u\} \succ m^* &
\end{aligned}$$

First suppose for some $\hat{u} \in U$ and $\hat{v} \in V$ we have $\langle \hat{u}, \hat{v} \rangle \geq l$ and let this be the pair with largest inner product. Consider the deferred acceptance algorithm for finding the man-optimal stable matching. First, some of the men corresponding to U will propose to the women corresponding to V and at least $m_{\hat{u}}$ will be accepted by $w_{\hat{v}}$. The remaining men corresponding to U will be accepted by dummy women. The dummy men will propose to the women corresponding to V but not all can be accepted. These rejected dummy men will propose to w^* who will accept one. Then when m^* proposes to w^* she will reject him, as will the women corresponding to V , so he will be matched with a dummy woman. Since m^* and w^* are not matched in the man optimal stable matching, (m^*, w^*) is not a pair in any stable matching.

Now suppose $\langle u, v \rangle < l$ for every $u \in U, v \in V$ and consider the deferred acceptance algorithm for finding the woman-optimal stable matching. First, the dummy women will propose to the men corresponding to U and will be accepted. Then every woman w_v will propose to the dummy men and be accepted. Since every man besides m^* is matched with a woman he prefers to w^* , when she proposes to them, she will be rejected, so she will pair with m^* . Since w^* receives her least preferred choice in the woman optimal stable matching, (m^*, w^*) is a pair in every stable matching. Figure 4.3 demonstrates each of these cases.

We can amend the vectors from Theorem 4.9 as follows so that they realize the changed preferences with the attribute model.

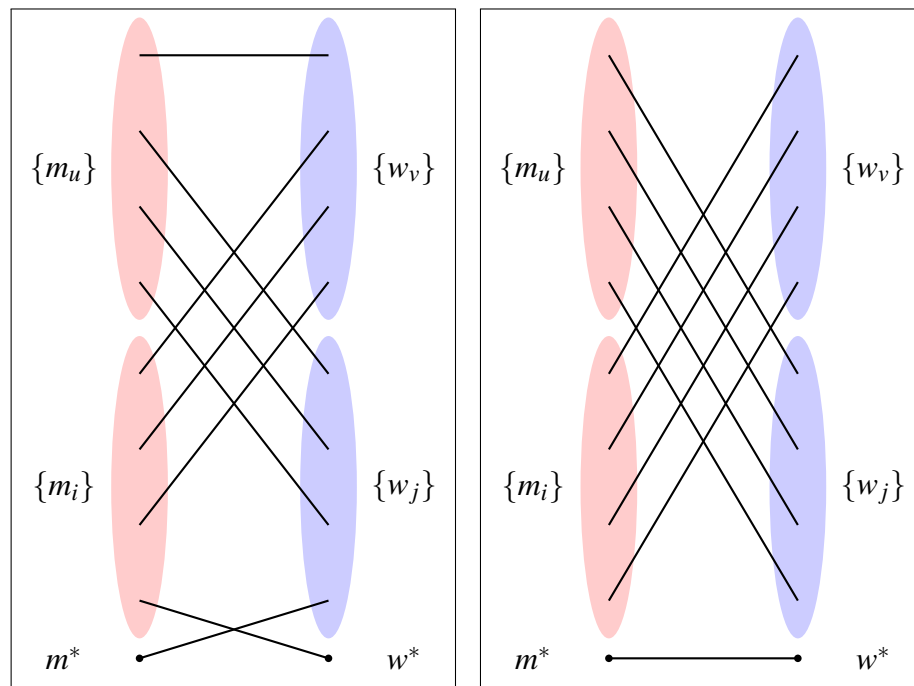
$$\begin{aligned}
m_u &: u^7 \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^6 \circ 0^6 \circ 0^6 \\
m_i &: 0^{7d} \circ 1^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 1^6 \circ 0^6 \\
m^* &: 0^{7d} \circ 1^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 0^6 \circ 1^6 \\
w_v &: v^7 \circ 0^{7l-1} \circ 1^{7l-1} \circ 0^6 \circ 1^6 \circ 1 \cdot 0^5 \\
w_j &: 0^{7d} \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^6 \circ \mathbf{1}^3 \circ \mathbf{0}^3 \circ 0^6 \\
w^* &: 0^{7d} \circ 1^{7l-1} \circ 0^{7l-1} \circ 1^3 \circ 0^3 \circ 1^4 \circ 0^2 \circ 1^2 \circ 0^4
\end{aligned}$$

□

We would like to point out that the results on the hardness for nondeterministic and co-nondeterministic algorithms do not apply to Merlin-Arthur (MA) algorithms, i.e. algorithms with access to both nondeterministic bits and randomness. Williams [88] gives fast MA algorithms for a number of SETH-hard problems, and the same techniques also yield a $O(dn)$ time MA algorithm for the verification of a stable matching in the boolean attribute model with d attributes. We can obtain MA algorithms with time $O(dn)$ for finding stable matchings and certifying that a pair is in at least one stable matching by first nondeterministically guessing a stable matching.

4.5 Other Succinct Preference Models

In this section, we provide subquadratic algorithms for other succinct preference models, single-peaked and geometric, which are motivated by economics.



(a) A representative stable matching when there is a pair (u, v) with $\langle u, v \rangle \geq l$ (b) A representative stable matching when $\langle u, v \rangle < l$ for every pair (u, v)

Figure 4.3. A representation of the reduction from maximum inner product to checking a stable pair such that a true maximum inner product instance maps to a false stable pair instance

4.5.1 One Dimensional Single-Peaked Preferences

Formally, we say the men's preferences over the women in a matching market are *single-peaked* if the women can be ordered as points along a line ($p(w_1) < p(w_2) < \dots < p(w_n)$) and for each man m there is a point $q(m)$ and a binary preference relation \succ_m such that if $p(w_i) \leq q(m)$ then $p(w_i) \succ_m p(w_j)$ for $j < i$ and if $p(w_i) \geq q(m)$ then $p(w_i) \succ_m p(w_j)$ for $j > i$. Essentially, each man prefers the women that are "closest" to his ideal point $q(m)$. One example of a preference relation for m would be the distance from $q(m)$. If the women's preferences are also single-peaked then we say the matching market has single peaked preferences. Since these preferences only consist of the p and q values and the preference relations for the participants, they can be represented succinctly as long as the relations require subquadratic space.

Verifying a Stable Matching for Single-Peaked Preferences

Here we demonstrate a subquadratic algorithm for verifying if a given matching is stable when the preferences of the matching market are single-peaked. We assume that the preference relations can be computed in constant time.

Theorem 4.11. *There is an algorithm to verify a stable matching in the single-peaked preference model in $O(n \log n)$ time.*

Proof. Let $p(m_i)$ be the point associated with man m_i , $q(m_i)$ be m_i 's preference point, and \succ_{m_i} be m_i 's preference relation. The women's points are denoted analogously. We assume that $p(m_i) < p(m_j)$ if and only if $i < j$ and the same for the women. Let μ be the given matching we are to check for stability.

First, for each man m , we compute the intervals along the line of women which includes all women m strictly prefers to $\mu(m)$. If this interval is empty, m is with his most preferred woman and cannot be involved in any blocking pairs so we can ignore him. For

all nonempty intervals each endpoint is $p(w)$ for some woman w . We also compute these intervals for the women. Note that for any man m and woman w , (m, w) is a blocking pair for μ if and only if m is in w 's interval and w is in m 's interval.

We will process each of the women in order from w_1 to w_n maintaining a balanced binary search tree of the men who prefer that woman to their partners. This will allow us to easily check if she prefers any of them by seeing if any elements in the tree are between the endpoints of her interval. Initially this tree is empty. When processing a woman w , we first add any man m whose interval begins with w to the search tree. Then we check to see if w prefers any men in the tree. If so, we know the matching is not stable. Otherwise, we remove any man m from the tree whose interval ends with w and proceed to the next woman. Algorithm 4.5 provides pseudocode for this algorithm.

Computing the intervals requires $O(n \log n)$. Since we only insert each man into the tree at most once, maintaining the tree requires $O(n \log n)$. The queries also require $O(\log n)$ for each woman so the total time is $O(n \log n)$. \square

Remarks on Finding a Stable Matching for Single-Peaked Preferences

The algorithm in [11] relies on the observation that there will always be a pair of participants who are each other's first choice with narcissistic single-peaked preferences. Thus a greedy approach where one such pair is selected and then removed works well. However, this is not the case when we remove the narcissistic assumption. In fact, as with the two-list case, Table 4.4 presents an example where no participant is matched with their top choice in the unique stable matching. Note that the preferences for the men and women are symmetric. The reader can verify that these preferences can be realized in the single-peaked preference model using the orderings $p(m_1) < p(m_2) < p(m_3) < p(m_4)$ and $p(w_1) < p(w_2) < p(w_3) < p(w_4)$ and that the unique stable matching is $\{(m_1, w_4), (m_2, w_2), (m_3, w_3), (m_4, w_1)\}$ where no participant receives their first choice.

Algorithm 4.5: Verify Single-Peaked Stable Matching

```

for each woman w do
  Create two empty lists w.begin and w.end.
  Use binary search to find the leftmost man m and rightmost man m' that
  w prefers to  $\mu(w)$  if any. (Otherwise remove w.)
  Let w.s =  $p(m)$  and w.t =  $p(m')$ .
for each man m do
  Use binary search to find the leftmost woman w and rightmost woman w'
  that m prefers to  $\mu(m)$  if any. (Otherwise ignore m.)
  Add m to w.begin and w'.end.
Initialize an empty balanced binary search tree T.
for i = 1 to n do
  for  $m \in w_i.begin$  do
     $T.insert(p(m))$ 
  if there are any points  $p(m)$  in T between  $w_i.s$  and  $w_i.t$  then
    return  $(m, w_i)$  is a blocking pair.
  for  $m \in w_i.end$  do
     $T.delete(p(m))$ 
return  $\mu$  is stable.

```

Table 4.4. Single-peaked preferences where no participant receives their top choice in the stable matching

Man	Preference List	Woman	Preference List
m_1	$w_3 \succ w_2 \succ w_4 \succ w_1$	w_1	$m_3 \succ m_2 \succ m_4 \succ m_1$
m_2	$w_3 \succ w_2 \succ w_4 \succ w_1$	w_2	$m_3 \succ m_2 \succ m_4 \succ m_1$
m_3	$w_4 \succ w_3 \succ w_2 \succ w_1$	w_3	$m_4 \succ m_3 \succ m_2 \succ m_1$
m_4	$w_2 \succ w_1 \succ w_3 \succ w_4$	w_4	$m_2 \succ m_1 \succ m_3 \succ m_4$

Also no greedy algorithm following the model inspired by [29] will succeed for single-peak preferences because the preferences in Table 4.2 can be realized in the single-peaked preference model using the orderings $p(m_1) < p(m_2) < p(m_3)$ and $p(w_1) < p(w_2) < p(w_3)$.

4.5.2 Geometric Preferences

We say the men’s preferences over the women in a matching market are *geometric* in d dimensions if each women w is defined by a *location* $p(w)$ and for each man m there is an *ideal* $q(m)$ such that m prefers woman w_1 to w_2 if and only if $\|p(m) - q(w_1)\|_2^2 < \|p(m) - q(w_2)\|_2^2$, i.e. $p(w_1)$ has smaller euclidean distance from the man’s ideal than $p(w_2)$. If the women’s preferences are also geometric we call the matching market *geometric*. We further call the preferences *narcissistic* if $p(x) = q(x)$ for every participant x . Our results for the attribute model extend to geometric preferences.

Note that one-dimensional geometric preferences are a special case of single-peaked preferences. As such, geometric preferences might be used to model preferences over political candidates who are given a score on several (linear) policy areas, e.g. protectionist vs. free trade and hawkish vs. dovish foreign policy.

Arkin et al. [7] also consider geometric preferences, but restrict themselves to the narcissistic case. Our algorithms do not require the preferences to be narcissistic, hence our model is more general. On the other hand, our lower bounds for large dimensions also apply to the narcissistic special case. While Arkin et al. take special care of different notions of stability in the presence of ties, we concentrate on weakly stable matchings. Although we restrict ourselves to the stable matching problem for the sake of presentation, all lower bounds and verification algorithms naturally extend to the stable roommate problem. Since all proofs in this section are closely related those for the attribute model, we restrict ourselves to proof sketches highlighting the main differences.

Theorem 4.1 extends immediately to the geometric case without any changes in the proof.

Corollary 4.3 (Geometric version of Theorem 4.1). *There is an algorithm to find a stable matching in the d -dimensional geometric model with at most a constant C distinct values in time $O(C^{2d}n(d + \log n))$.*

For the verification of a stable matching with real-valued vectors we use a standard lifting argument.

Corollary 4.4 (Geometric version of Theorem 4.4). *There is an algorithm to verify a stable matching in the d -dimensional geometric model with real-valued locations and ideals in time $\tilde{O}(n^{2-1/2(d+1)})$*

Proof. Let $q \in \mathbb{R}^d$ be an ideal and let $a, b \in \mathbb{R}^d$ be two locations. Define $q', a', b' \in \mathbb{R}^{d+1}$ as $q' = (q_1, \dots, q_d, -1/2)$, $a' = (a_1, \dots, a_d, \sum_{i=1}^d a_i^2)$ and $b' = (b_1, \dots, b_d, \sum_{i=1}^d b_i^2)$.

We have $\langle a', q \rangle = 1/2 \sum_{i=1}^d q_i - 1/2 \|q - a\|_2^2$. Hence we get $\|q - a\|_2^2 < \|q - b\|_2^2$ if and only if $\langle q', a' \rangle > \langle q', b' \rangle$, so we can reduce the stable matching problem in the d -dimensional geometric model to the $d + 1$ -attribute model. \square

For the boolean case, we can adjust the proof of Theorem 4.6 by using a threshold of parities instead of a threshold of conjunctions. The degree of the resulting polynomial remains the same.

Corollary 4.5 (Geometric version of Theorem 4.6). *In the geometric model with n men and women, with locations and ideals in $\{0, 1\}^d$ with $d = c \log n$, there is a randomized algorithm to decide if a given matching is stable in time $\tilde{O}(n^{2-1/O(c \log^2(c))})$ with error probability at most $1/3$.*

For lower bounds we reduce from the minimum Hamming distance problem which is SETH-hard with the same parameters as the maximum inner product problem

[6]. The Hamming distance of two boolean vectors is exactly their squared euclidean distance, hence a matching market where the preferences are defined by Hamming distances is geometric.

Definition 4.5. *For any d and input l , the minimum Hamming distance problem is to decide if two input sets $U, V \subseteq \{0, 1\}^d$ with $|U| = |V| = n$ have a pair $u \in U, v \in V$ such that $\|u - v\|_2^2 < l$.*

Lemma 4.10 ([6]). *Assuming SETH, for any $\varepsilon > 0$, there is a c such that solving the minimum Hamming distance problem on $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

For the hardness of finding a stable matching, the construction from Theorem 4.7 works without adjustments.

Corollary 4.6 (Geometric version of Theorem 4.7). *Assuming SETH, for any $\varepsilon > 0$, there is a c such that finding a stable matching in the (boolean) d -dimensional geometric model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

For the hardness of verifying a stable matching, the construction is as follows.

Corollary 4.7 (Geometric version of Theorem 4.8). *Assuming SETH, for any $\varepsilon > 0$, there is a c such that verifying a stable matching in the (boolean) d -dimensional geometric model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. Let $U, V \subseteq \{0, 1\}^d$ be the inputs to the minimum Hamming distance problem and let l be the threshold.

For every $u \in U$, define a real man m_u with both ideal and location as $u \circ 0^l$ and a dummy woman w'_u with ideal and location $u \circ 1^l$. Symmetrically for $v \in V$ define w_v with $v \circ 0^l$ and m'_v with $v \circ 1^l$. The matching (m_u, w'_u) for all $u \in U$ and (w_v, m'_v) for all $v \in V$ is stable if and only if there is no pair u, v with Hamming distance less than l . \square

The hardness results for checking a stable pair also translate to the geometric model. In particular, since both variants of the proof extend to the geometric model we have the same consequences for nondeterministic algorithms.

Corollary 4.8 (Geometric version of Theorem 4.9). *Assuming SETH, for any $\varepsilon > 0$, there is a c such that determining whether a given pair is part of any or all stable matchings in the (boolean) d -dimensional geometric model with $d = c \log n$ dimensions requires time $\Omega(n^{2-\varepsilon})$.*

Proof. We again reduce from the minimum Hamming distance problem. We assume without loss of generality that d is even and the threshold l is exactly $d/2 + 1$, i.e. the instance is true if and only if there are vectors u, v with Hamming distance at most $d/2$. We can reduce to this case from any other threshold by padding the vectors.

We use the same preference orders as in the d -attribute model. The following narcissistic instance realizes the preference order from Theorem 4.9. For a vector $u \in \{0, 1\}^d$, \bar{u} denotes its component-wise complement.

$$\begin{aligned}
 m_u &: (u \circ \bar{u} \circ u \circ \bar{u})^3 \circ 000000000 \\
 m_i &: 0^{12d} \quad \circ 100000000 \\
 m^* &: 0^{12d} \quad \circ 001111111 \\
 w_v &: (v \circ \bar{v} \circ v \circ \bar{v})^3 \circ 000000000 \\
 w_j &: (0^{2d} \circ 1^{2d})^3 \quad \circ 010000000 \\
 w^* &: (0^{2d} \circ 1^{2d})^3 \quad \circ 101110000
 \end{aligned}$$

Likewise the preference orders for Theorem 4.10 are achieved by the following

vectors.

$$m_u : (u \circ \bar{u} \circ u \circ \bar{u})^3 \circ 000000000000$$

$$m_i : 0^{12d} \quad \circ 100000000000$$

$$m^* : 0^{12d} \quad \circ 001111111100$$

$$w_v : (v \circ \bar{v} \circ v \circ \bar{v})^3 \circ 000000000000$$

$$w_j : (0^{2d} \circ 1^{2d})^3 \quad \circ 01000000011$$

$$w^* : (0^{2d} \circ 1^{2d})^3 \quad \circ 10111000000$$

□

4.5.3 Strategic Behavior

With geometric and single-peaked preferences, we assume that the participants are not allowed to misrepresent their location points. Rather they may only misrepresent their preference ideal. As such, the results of this section do not apply when preferences are narcissistic.

Theorem 4.12. *There is no strategy proof algorithm to find a stable matching in the geometric preference model.*

Proof. We consider one-dimensional geometric preferences. Let the preference points and ideals be as given in Table 4.5 which yield the depicted preference lists. As in the proof for 4.3, there are two stable matchings: $\{(m_1, w_3), (m_2, w_1), (m_3, w_2)\}$, the man-optimal matching, and $\{(m_1, w_3), (m_2, w_2), (m_3, w_1)\}$, the woman-optimal matching. However, if w_2 changes her ideal to $5/3$ then her preference list is $m_2 \succ m_1 \succ m_3$. Now

Table 4.5. Geometric preferences that can be manipulated

Man	Location (p)	Ideal (q)	Woman	Location (p)	Ideal (q)
m_1	1	$7/3$	w_1	1	3
m_2	2	1	w_2	2	$7/3$
m_3	3	$5/3$	w_3	3	3

Man	Preference List	Woman	Preference List
m_1	$w_2 \succ w_3 \succ w_1$	w_1	$m_3 \succ m_2 \succ m_1$
m_2	$w_1 \succ w_2 \succ w_3$	w_2	$m_2 \succ m_3 \succ m_1$
m_3	$w_2 \succ w_1 \succ w_3$	w_3	$m_3 \succ m_2 \succ m_1$

there is a unique stable matching which is $\{(m_1, w_3), (m_2, w_2), (m_3, w_1)\}$, the woman-optimal stable matching from the original preferences. Therefore, any mechanism that does not always output the woman optimal stable matching can be manipulated by the women to their advantage. Similarly, any mechanism that does not always output the man-optimal matching could be manipulated by the men in some instances. Thus there is no strategy-proof mechanism for geometric preferences. \square

Since one-dimensional geometric preferences are a special case of single-peaked preferences the following corollary results directly from Theorem 4.12.

Corollary 4.9. *There is no strategy proof algorithm to find a stable matching in the single-peaked preference model.*

4.6 Conclusion and Open Problems

We give subquadratic algorithms for finding and verifying stable matchings in the d -attribute model and d -list model. We also show that, assuming SETH, one can only hope to find such algorithms if the number of attributes d is bounded by $O(\log n)$.

For a number of cases there is a gap between the conditional lower bound and the upper bound. Our algorithms with real attributes and weights are only subquadratic if the

dimension is constant. Even for small constants our algorithm to find a stable matching is not tight, as it is not subquadratic for any $d = O(\log n)$. The techniques we use when the attributes and weights are small constants do not readily apply to the more general case.

There is also a gap between the time complexity of our algorithms for finding a stable matching and verifying a stable matching. It would be interesting to either close or explain this gap. On the one hand, subquadratic algorithms for finding a stable matching would demonstrate that the attribute and list models are computationally simpler than the general preference model. On the other hand, proving that there are no subquadratic algorithms would show a distinction between the problems of finding and verifying a stable matching in these settings which does not exist for the general preference model. Currently, we do not have a subquadratic algorithm for finding a stable matching even in the 2-list case, while we have an optimal algorithm for verifying a stable matching for d lists. This 2-list case seems to be a good starting place for further research.

Additionally it is worth considering succinct preference models for other computational problems that involve preferences to see if we can also develop improved algorithms for these problems. For example, the Top Trading Cycles algorithm [82] can be made to run in subquadratic time for d -attribute preferences (when d is constant) using the ray shooting techniques applied in this chapter to find participants' top choices.

4.7 Acknowledgements

Chapter 4, in part, is a reprint of the material as it appears in International Computer Science Symposium in Russia 2016, pages 298–308. Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. “Subquadratic Algorithms for Succinct Stable Matching.” Springer, 2016. [68] The dissertation author was the primary investigator and author of this paper.

Chapter 4, in part, is currently being prepared for submission for publication

of the material. Künnemann, Marvin, Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. “Subquadratic Algorithms for Succinct Stable Matching.” arXiv preprint arXiv:1510.06452 (2016). [59] The dissertation author was the primary investigator and author of this material.

Bibliography

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Quadratic-time hardness of LCS and other sequence similarity measures. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 2015.
- [2] David J Abraham, Katarína Cechlárová, David F Manlove, and Kurt Mehlhorn. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*, pages 1163–1175. Springer, 2005.
- [3] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. *SIAM J. Comput.*, 40(1):92–106, 2011.
- [4] Pankaj K Agarwal, Lars Arge, Jeff Erickson, Paolo G Franciosa, and Jeffry Scott Vitter. Efficient searching with linear constraints. In *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 169–178. ACM, 1998.
- [5] Pankaj K. Agarwal and Jeff Erickson. Geometric range searching and its relatives. *Contemporary Mathematics*, 223:1–56, 1999.
- [6] Josh Alman and Ryan Williams. Probabilistic polynomials and hamming nearest neighbors. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 2015.
- [7] Esther M. Arkin, Sang Won Bae, Alon Efrat, Kazuya Okamoto, Joseph S.B. Mitchell, and Valentin Polishchuk. Geometric stable roommates. *Information Processing Letters*, 109(4):219–224, 2009.
- [8] Kenneth J Arrow. A difficulty in the concept of social welfare. *The Journal of Political Economy*, pages 328–346, 1950.
- [9] Itai Ashlagi and Flip Klijn. Manipulability in matching markets: conflict and coincidence of interests. *Social Choice and Welfare*, 39(1):23–33, 2012.
- [10] Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the Forty-Seventh*

Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pages 51–58, 2015.

- [11] John Bartholdi and Michael A. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.
- [12] Nayantara Bhatnagar, Sam Greenberg, and Dana Randall. Sampling stable marriages: why spouse-swapping won't work. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1223–1232. Society for Industrial and Applied Mathematics, 2008.
- [13] Péter Biró and Gethin Norman. Analysis of stochastic matching markets. *International Journal of Game Theory*, 42(4):1021–1040, 2013.
- [14] Duncan Black. The decisions of a committee using a special majority. *Econometrica: Journal of the Econometric Society*, pages 245–261, 1948.
- [15] Anna Bogomolnaia and Jean-François Laslier. Euclidean preferences. *Journal of Mathematical Economics*, 43(2):87–98, 2007.
- [16] James W. Boudreau. Preference structure and random paths to stability in matching markets. *Economics Bulletin*, 3(67):1–12, 2008.
- [17] James W Boudreau. A note on the efficiency and fairness of decentralized matching. *Operations Research Letters*, 39(4):231–233, 2011.
- [18] James W. Boudreau. An exploration into why some matchings are more likely than others. In *MATCH-UP 2012: the Second International Workshop on Matching Under Preferences*, page 39, 2012.
- [19] James W Boudreau and Vicki Knoblauch. Marriage matching and intercorrelation of preferences. *Journal of Public Economic Theory*, 12(3):587–602, 2010.
- [20] Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless seth fails. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 661–670. IEEE, 2014.
- [21] Guido Caldarelli and Andrea Capocci. Beauty and distance in the stable marriage problem. *Physica A: Statistical Mechanics and its Applications*, 300(1):325–331, 2001.
- [22] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 261–270. ACM, 2016.

- [23] Onur B Celik and Vicki Knoblauch. Marriage matching with correlated preferences. *Economics Working Papers*, page 200716, 2007.
- [24] Prasad Chebolu, Leslie Ann Goldberg, and Russell Martin. The complexity of approximately counting stable matchings. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 81–94. Springer, 2010.
- [25] Christine T Cheng. On the stable matchings that can be reached when the agents go marching in one by one. *Accepted to SIAM Journal on Discrete Mathematics*, 2016.
- [26] Kim-Sau Chung. On the existence of stable roommate matchings. *Games and economic behavior*, 33(2):206–230, 2000.
- [27] Peter Coles and Ran Shorrer. Optimal truncation in matching markets. *Games and Economic Behavior*, 87:591–615, 2014.
- [28] John Dabney and Brian C. Dean. Adaptive stable marriage algorithms. In *Proceedings of the 48th Annual Southeast Regional Conference*, page 35. ACM, 2010.
- [29] Sashka Davis and Russell Impagliazzo. Models of greedy algorithms for graph problems. *Algorithmica*, 54(3):269–317, 2009.
- [30] David P. Dobkin and David G. Kirkpatrick. A linear algorithm for determining the separation of convex polyhedra. *Journal of Algorithms*, 6(3):381–392, 1985.
- [31] Federico Echenique and Leeat Yariv. An experimental study of decentralized matching. *Economics Working Papers*, 2012.
- [32] Tomás Feder. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences*, 45(2):233–284, 1992.
- [33] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [34] David Gale and Marilda Sotomayor. Ms. machiavelli and the stable matching problem. *The American Mathematical Monthly*, 92(4):261–268, 1985.
- [35] David Gale and Marilda Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3):223–232, 1985.
- [36] Yannai A. Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1003–1017. SIAM, 2015.

- [37] Dan Gusfield. Three fast algorithms for four problems in stable marriage. *SIAM Journal on Computing*, 16(1):111–128, 1987.
- [38] Dan Gusfield and Robert W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. Foundations of Computing Series. Mit Press, 1989.
- [39] John Hershberger and Subhash Suri. A pedestrian approach to ray shooting: Shoot a ray, take a walk. *Journal of Algorithms*, 18(3):403–431, 1995.
- [40] Martin Hoefer. Local matching dynamics in social networks. *Information and Computation*, 222:20–35, 2013.
- [41] Martin Hoefer and Lisa Wagner. Matching dynamics with constraints. In *International Conference on Web and Internet Economics*, pages 161–174. Springer, 2014.
- [42] Moshe Hoffman, Daniel Moeller, and Ramamohan Paturi. Jealousy graphs: Structure and complexity of decentralized stable matching. In *International Conference on Web and Internet Economics*, pages 263–276. Springer, 2013.
- [43] Chien-Chung Huang. Cheating by men in the gale-shapley stable matching algorithm. In *European Symposium on Algorithms*, pages 418–431. Springer, 2006.
- [44] Nicole Immorlica and Mohammad Mahdian. Marriage, honesty, and stability. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 53–62. Society for Industrial and Applied Mathematics, 2005.
- [45] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367 – 375, 2001.
- [46] Russell Impagliazzo, Ramamohan Paturi, and Stefan Schneider. A satisfiability algorithm for sparse depth two threshold circuits. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 479–488. IEEE, 2013.
- [47] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63:512–530, 2001.
- [48] Robert W Irving. An efficient algorithm for the stable roommates problem. *Journal of Algorithms*, 6(4):577–595, 1985.
- [49] Robert W Irving. Stable marriage and indifference. *Discrete Applied Mathematics*, 48(3):261–272, 1994.
- [50] Robert W Irving and Paul Leather. The complexity of counting stable marriages. *SIAM Journal on Computing*, 15(3):655–667, 1986.

- [51] Robert W Irving, Paul Leather, and Dan Gusfield. An efficient algorithm for the optimal stable marriage. *Journal of the ACM (JACM)*, 34(3):532–543, 1987.
- [52] Kazuo Iwama and Shuichi Miyazaki. A survey of the stable marriage problem and its variants. In *Informatics Education and Research for Knowledge-Circulating Society, 2008. ICKS 2008. International Conference on*, pages 131–136. IEEE, 2008.
- [53] Kazuo Iwama, Shuichi Miyazaki, Yasufumi Morita, and David Manlove. Stable marriage with incomplete lists and ties. In *International Colloquium on Automata, Languages, and Programming*, pages 443–452. Springer, 1999.
- [54] Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna Paluch. Strongly stable matchings in time $o(nm)$ and extension to the hospitals-residents problem. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 222–233. Springer, 2004.
- [55] Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science*, 127(2):255–267, 1994.
- [56] Donald E. Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*, volume 10. Amer Mathematical Society, 1997.
- [57] Hirotatsu Kobayashi and Tomomi Matsui. Cheating strategies for the gale-shapley algorithm with complete preference lists. *Algorithmica*, 58(1):151–169, 2010.
- [58] Fuhito Kojima and Parag A Pathak. Incentives and stability in large two-sided matching markets. *The American Economic Review*, 99(3):608–627, 2009.
- [59] Marvin Künnemann, Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. Subquadratic algorithms for succinct stable matching. *arXiv preprint arXiv:1510.06452*, 2016.
- [60] Dai Tri Man Lê, Stephen A Cook, and Yuli Ye. A formal theory for the complexity class associated with the stable marriage problem. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 12. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2011.
- [61] Jinpeng Ma. Stable matchings and rematching-proof equilibria in a two-sided matching market. *Journal of Economic Theory*, 66(2):352–369, 1995.
- [62] David F Manlove. Stable marriage with ties and unacceptable partners. *Technical Report TR-1999-29 University of Glasgow, Department of Computing Science*, 1999.

- [63] David F Manlove, Robert W Irving, Kazuo Iwama, Shuichi Miyazaki, and Yasufumi Morita. Hard variants of stable marriage. *Theoretical Computer Science*, 276(1): 261–279, 2002.
- [64] Jiří Matoušek. Efficient partition trees. *Discrete & Computational Geometry*, 8(1):315–334, 1992.
- [65] Jiří Matoušek and Otfried Schwarzkopf. Linear optimization queries. In *Proceedings of the eighth annual symposium on Computational geometry*, pages 16–25. ACM, 1992.
- [66] Tomomi Matsui. Algorithmic aspects of equilibria of stable marriage model with complete preference lists. In *Operations Research Proceedings 2010*, pages 47–52. Springer, 2011.
- [67] Ernst W Mayr and Ashok Subramanian. The complexity of circuit value and network stability. In *Structure in Complexity Theory Conference, 1989. Proceedings., Fourth Annual*, pages 114–123. IEEE, 1989.
- [68] Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. Subquadratic algorithms for succinct stable matching. In *International Computer Science Symposium in Russia*, pages 294–308. Springer, 2016.
- [69] Susan Mongell and Alvin E Roth. Sorority rush as a two-sided matching mechanism. *The American Economic Review*, pages 441–464, 1991.
- [70] Cheng Ng and Daniel S. Hirschberg. Lower bounds for the stable marriage problem and its variants. *SIAM Journal on Computing*, 19(1):71–77, 1990.
- [71] Joana Pais, Agnes Pinter, and Robert F Veszteg. Decentralized matching markets: a laboratory experiment. *Economics Working Papers*, 2012.
- [72] Mihai Patrascu and Ryan Williams. On the possibility of faster SAT algorithms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 1065–1075, 2010.
- [73] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes*, 41(4):333–338, 1987.
- [74] Alvin E. Roth. The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4):617–628, 1982.
- [75] Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *The Journal of Political Economy*, pages 991–1016, 1984.

- [76] Alvin E Roth. Two-sided matching with incomplete information about others' preferences. *Games and Economic Behavior*, 1(2):191–209, 1989.
- [77] Alvin E Roth and Uriel G Rothblum. Truncation strategies in matching markets in search of advice for participants. *Econometrica*, 67(1):21–43, 1999.
- [78] Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-sided Matching: A Study in Game - Theoretic Modeling and Analysis*. Econometric Society Monographs. Cambridge University, 1990.
- [79] Alvin E. Roth and John H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, nov 1990.
- [80] Alvin E Roth and John H Vande Vate. Incentives in two-sided matching with random stable mechanisms. *Economic theory*, 1(1):31–44, 1991.
- [81] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136(1):341–378, 2007.
- [82] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.
- [83] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82. ACM, 1987.
- [84] Ashok Subramanian. A new approach to stable matching problems. *SIAM Journal on Computing*, 23(4):671–700, 1994.
- [85] Chung-Piaw Teo, Jay Sethuraman, and Wee-Peng Tan. Gale-shapley stable marriage problem revisited: Strategic issues and applications. *Management Science*, 47(9):1252–1267, 2001.
- [86] Ryan Williams. A new algorithm for optimal constraint satisfaction and its implications. In *Automata, Languages and Programming*, pages 1227–1237. Springer, 2004.
- [87] Ryan Williams. Faster all-pairs shortest paths via circuit complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 664–673. ACM, 2014.
- [88] Ryan Williams. Strong ETH breaks with merlin and arthur: Short non-interactive proofs of batch evaluation. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, pages 2:1–2:17, 2016.