Jealousy Graphs: Structure and Complexity of Decentralized Stable Matching

Moshe Hoffman, Daniel Moeller, and Ramamohan Paturi

University of California, San Diego
`http://www.ucsd.edu`

Abstract. The stable matching problem has many applications to real world markets and efficient centralized algorithms are known. However, little is known about the decentralized case. Several natural randomized algorithmic models for this setting have been proposed but they have worst case exponential time in expectation. We present 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.

Keywords: decentralized stable matching, market algorithms

1 Introduction

The stable matching problem and its variants have been widely studied due to real world market applications, such as assigning residents to hospitals, women to sororities, and students to public schools [1,2,3]. In a seminal paper, Gale and Shapley first proposed an algorithm to find a stable matching in the basic two-sided (bipartite) version [4]. Others have subsequently investigated the structure of the set of stable matchings [1]. However, most prior work involves centralized algorithms to find stable matchings 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 matchings 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 processes will find a stable matching and provides experimental support [5]. 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 [6]. At each step these algorithms match two participants who form a blocking pair (who prefer to be matched with each other over its partner) of the current matching. However, they present no expected time complexity. Ackerman 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 [7].

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 [7]. However, 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 paper 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, whereas our structures relate to individual stable matchings and the distributed process by which these stable matchings are achieved [1].

With a decomposition, we associate a size and depth. Our main result, Theorem 24, 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 24 provides theoretical justification for the simulated results of Boudreau [8]. We also conjecture that these structures will 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 [9,5,10].

In the remainder of this section we formalize our model and present the basic concepts. In section 2, we present some useful structural properties of the jealousy graph and decomposition. In section 3, we have our convergence result. Section 4 contains the application of our work to [8], and section 5 is our conclusion.

1.1 Basic Stable Matching Concepts

We start with the basic definitions of matching markets and stable matchings. Those familiar with the matching literature will notice that we restrict preferences to be complete and strict.

Definition 1 *(S, P) is a matching market if $S = M \bigcup W$ 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, and for $s \in W$, $\succ_s$ is a total order over M.*

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

Definition 2 *A matching on the set S is a function $\mu : S \to S$ such that $\forall s \in S$, $\mu(\mu(s)) = s$, $s \in M \Rightarrow \mu(s) \in W \bigcup \{s\}$ and $s \in W \Rightarrow \mu(s) \in M \bigcup \{s\}$.*

We say that a participant $s \in S$ is unmatched by a matching μ if $\mu(s) = s$. We also assume that all participants prefer to be matched to anyone than to be unmatched. 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 3 *A matching on the set S is a perfect matching if $\mu(s) \neq s$ for all $s \in S$.*

Given a matching, if there were a man and a woman who each preferred the other to their partner, this causes 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 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 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).*

The following three concepts will be useful since we will deal with subsets of the matching market.

Definition 6 *A balanced subset of a matching market (S, P), $S = M \bigcup W$, is a subset $S' \subseteq S$ such that $|S' \bigcap M| = |S' \bigcap W|$.*

Definition 7 *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 8 *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 $\mu' \upharpoonright_{S'} = \mu \upharpoonright_{S'}$.*

1.2 Better Response Algorithm

The class of algorithms introduced by Roth and Vande Vate [6] 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 9 *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 [7], 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.

1.3 Jealousy Graph and Related Definitions

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 will provide bounds on convergence based on this simpler structure.

Definition 10 *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* $E = \{(u_{\{x,y\}}, v_{\{x',y'\}}) | u_{\{x,y\}}, v_{\{x',y'\}} \in V$, *and either* $x \succ_{y'} x'$ *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 11 *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 will be 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 will 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 will 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 $Down(A) = \{B | B \preceq A\}$. When the elements of Π are sets themselves, as in the case of decompositions, we will denote 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 12 *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 will 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 13 *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.

2 Structural Results

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 14 *Given any directed graph G with n vertices, there is a set $S = \{m_i, w_i | i = 1, 2, \ldots, n\}$ and preferences $P = \{\succ_{m_i}, \succ_{w_i} | 1 \leq i \leq n\}$ 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 $S = \{m_i, w_i | i = 1, 2, \ldots, n\}$. Arbitrarily index the vertices in V as $v_1, v_2, \ldots, v_n$. We will design our preferences such that if $\mu(m_i) = w_i$ for $i = 1, 2, \ldots, n$, μ is a stable matching and v_i is the vertex corresponding to $\{m_i, w_i\}$. Now define $P = \{\succ_{m_i}, \succ_{w_i} | 1 \leq i \leq n\}$ 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 $\{w_j | w_j \succ_{m_i} w_i\}$ and $\{w_j | w_i \succ_{m_i} w_j\}$ 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 $V' = v'_1, v'_2, \ldots, v'_n$. We will let v'_i correspond to $\{m_i, w_i\}$ for $i = 1, 2, \ldots, n$. Since all women are matched with their top preference by μ, the women will not be 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.

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 15 *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 $v_1, v_2, \ldots, v_k$ 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 $v_{\{m,w\}} \in A_1$.

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 16 *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 15, $s \in \mathbf{D}(A)$. This is a contradiction.

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 17 *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 $\mu(A \bigcap B) = \mu'(A \bigcap B) = A \bigcap B$.*

Proof. Suppose there is $x \in A \bigcap B$ but $\mu(x) \in A - B$. By Lemma 16, $\mu'(x) \succ_x \mu(x)$. In that case, $\mu'(x) \in A \bigcap B$, also by Lemma 16. 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 16, $\mu(\mu'(x)) \in A \bigcap B$. Continuing in this manner gives an infinite sequence $x, \mu'(x), \mu(\mu'(x)), \mu'(\mu(\mu'(x))), \ldots \in A \bigcap B$. These elements are distinct elements since $\mu(x) \neq \mu'(x)$ and both μ and μ' are bijective. This is a contradiction since $A \bigcap B$ is finite. Therefore $\mu(A \bigcap B) = \mu'(A \bigcap B) = A \bigcap B$.

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 18 *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 16 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 17, for each stable component B of ρ_μ, $\mu(B-\mathbf{D}(A)) = B-\mathbf{D}(A)$. The set $\{B-\mathbf{D}(A)|B \in \rho_\mu\}$ 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.

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 24 uses this to bound the time it takes for all components of the decomposition to reach a stable matching.

For this section we will assume (S, P) is a matching market of size n, μ is a stable matching on S, and $(\Pi, \preceq)$ 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 19 *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'.

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 20 *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 19 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.

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 21 *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 $\mu_0 \restriction_A$ contains m of the pairs in μ where $0 \leq m < c$. Let $\mu_0, \mu_1, \ldots, \mu_t$ 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 $\exists j, 0 < j \leq t, \mu_j \restriction_A$ contains at least $m+1$ of the pairs in μ is at least $\frac{1}{c^4}$.

Proof. Assume $\mu_0, \mu_1, \ldots, \mu_t$ 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, $\mu_{j-1} \restriction_A$ has m pairs in μ. By lemma 20 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 $\mu_j \restriction_A$ 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', $\mu_i \restriction_A$ 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, $\mu_i \restriction_A$ 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 $\mu_{l-1} \restriction_A$ must have $m-1$ pairs that are in μ. By lemma 20 there is at least $\frac{1}{c^2}$ probability that $\mu_l \restriction_A$ 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, $l < j \leq t$, $\mu_j \restriction_A$ has $m+1$ pairs that are in μ. This gives us a total probability of at least $\frac{1}{c^4}$ that $\mu_j \restriction_A$ has $m+1$ pairs that are in μ for some j, $0 < j \leq t$.

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

Lemma 22 *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 $\{0, 1, \ldots, c\}$ 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 21, 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 $i = 0, 1, \ldots, n$.

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 23 *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 19 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.

Finally we will 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 24 (Convergence) *Suppose μ is a stable matching. Suppose the depth of $(\Pi, \preceq)$ 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 $(\Pi, \preceq)$. Recall that a stable component A is on level j if $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 $\mathcal{D}$, $n_1 + \ldots + n_d \leq n$. We will 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 23, we know it will take cn^{2c} steps in expectation to reach a locally perfect matching on A. Also, by lemma 21 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 23, we know it will take cn^{2c} steps in expectation to reach a locally perfect matching on A. By Lemma 22, 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

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

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 18, 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 23 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 $T(0) =\leq \sum_{i=1}^{d} n_i n^2 = n^3$.

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

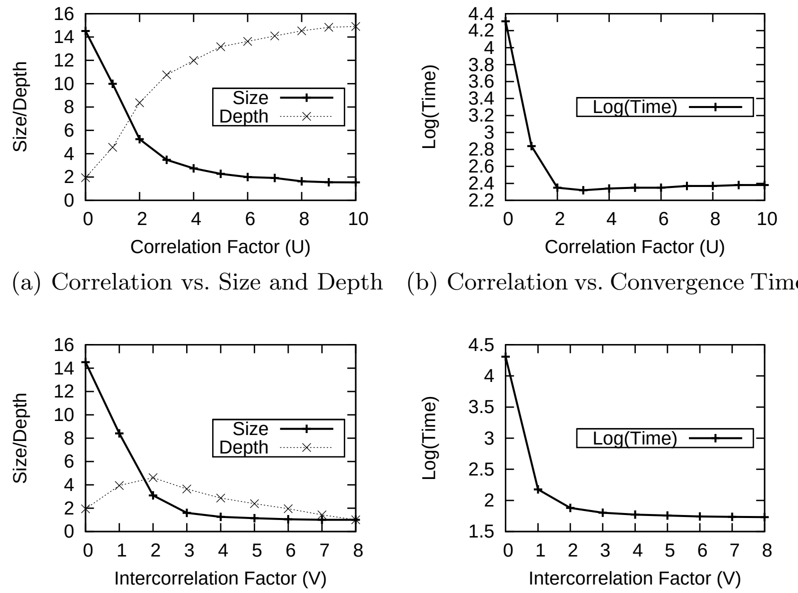

(a) Correlation vs. Size and Depth (b) Correlation vs. Convergence Time

(c) Intercorrelation vs. Size and Depth (d) Intercorrelation vs. Convergence Time

Fig. 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.

small size components. Partially correlated preferences are often used by modelers [8,11] 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 Ackemann et al. Intercorrelation exists when the preferences of the men relate to the preferences of the women. See [12] for examples of markets with intercorrelation. Boudreau showed that more correlation and intercorrelation lead to faster convergence of the better response algorithm [8]. We provide similar plots in Figures 1(b) and 1(d). Theorem 24 provides theoretical justification for these simulated results.

As described in [13,11], 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 1(a). 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 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 1(b) shows the log of the average convergence time over 100 trials for each of the 100 correlated preferences generated.

As in [12], 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. $1(c)$ and $1(d)$ are generated in the same manner as $1(a)$ and $1(b)$, respectively. These plots show that as preferences become more intercorrelated, the size and depth of the decompositions decrease. As Theorem 24 explains, this decreases the convergence time of the better response algorithm as intercorrelation increases.

5 Conclusion

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.

Acknowledgments This research was partially supported by the National Science Foundation under Grant No. 0905645, by the Army Research Office grant number W911NF-11-1-0363, and by grant RFP-1211 from the Foundational Questions in Evolutionary Biology Fund. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The final publication is available at Springer via `http://dx.doi.org/10.1007/978-3-319-34171-2_21`.

References

1. Gusfield, D., Irving, R.: The Stable Marriage Problem: Structure and Algorithms. Foundations of Computing Series. Mit Press (1989)
2. Roth, A., Sotomayor, M.: Two-sided Matching: A Study in Game - Theoretic Modeling and Analysis. Econometric Society Monographs. Cambridge University (1990)
3. Knuth, D.E.: Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms. Volume 10. Amer Mathematical Society (1997)
4. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. The American Mathematical Monthly **69**(1) (1962) 9–15
5. Echenique, F., Yariv, L.: An experimental study of decentralized matching (2011) Working paper, Caltech.
6. Roth, A.E., Vande Vate, J.H.: Random paths to stability in two-sided matching. Econometrica **58**(6) (nov 1990) 1475–1480
7. Ackermann, H., Goldberg, P.W., Mirrokni, V.S., Röglin, H., Vöcking, B.: Uncoordinated two-sided matching markets. SIAM J. Comput. **40**(1) (2011) 92–106
8. Boudreau, J.W.: Preference structure and random paths to stability in matching markets. Economics Bulletin **3**(67) (2008) 1–12
9. Boudreau, J.W.: An exploration into why some matchings are more likely than others. In: MATCH-UP 2012: the Second International Workshop on Matching Under Preferences. (2012) 39
10. Pais, J., Pinter, A., Veszteg, R.F.: Decentralized matching markets: a laboratory experiment. (2012)
11. Celik, O.B., Knoblauch, V.: Marriage matching with correlated preferences. Economics Working Papers (2007) 200716
12. Boudreau, J.W., Knoblauch, V.: Marriage matching and intercorrelation of preferences. Journal of Public Economic Theory **12**(3) (2010) 587–602
13. Caldarelli, G., Capocci, A.: Beauty and distance in the stable marriage problem. Physica A: Statistical Mechanics and its Applications **300**(1) (2001) 325–331