

Locally Stable Marriage with Strict Preferences^{*}

Martin Hoefer¹ and Lisa Wagner²

¹ Max-Planck-Institut für Informatik and Saarland University, Germany
mhoefer@mpi-inf.mpg.de

² Dept. of Computer Science, RWTH Aachen University, Germany
lwagner@cs.rwth-aachen.de

Abstract. We study two-sided matching markets with locality of information and control. Each male (female) agent has an arbitrary strict preference list over all female (male) agents. In addition, each agent is a node in a fixed network. Agents learn about possible partners dynamically based on their current network neighborhood. We consider convergence of dynamics to locally stable matchings that are stable with respect to their imposed information structure in the network. While existence of such states is guaranteed, we show that reachability becomes NP-hard to decide. This holds even when the network exists only among one side. In contrast, if only one side has no network and agents remember a previous match every round, reachability is guaranteed and random dynamics converge with probability 1. We characterize this positive result in various ways. For instance, it holds for random memory and for memory with the most recent partner, but not for memory with the best partner. Also, it is crucial which partition of the agents has memory. Finally, we conclude with results on approximating maximum locally stable matchings.

1 Introduction

Matching problems form the basis of many assignment and allocation tasks encountered in computer science, operations research, and economics. A prominent and popular approach in all these areas is *stable matching*, as it captures aspects like distributed control and rationality of participants that arise in many assignment problems today. A variety of allocation problems in markets can be analyzed within the context of two-sided stable matching, e.g., the assignment of jobs to workers [2, 5], organs to patients [18], or general buyers to sellers. In addition, stable marriage problems have been successfully used to study distributed resource allocation problems in networks [9].

In this paper, we consider a game-theoretic model for matching with distributed control and information. Agents are rational agents embedded in a (social) network and strive to find a partner for a joint relationship or activity,

^{*} Supported by DFG Cluster of Excellence MMCI and grant Ho 3831/3-1.

An extended full version of this paper can be found at <http://arxiv.org/abs/1207.1265>

e.g., to do sports, write a research paper, exchange data etc. Such problems are of central interest in economics and sociology, and they act as fundamental coordination tasks in distributed computer networks. Our model extends the stable marriage problem, in which we have sets U and W of men and women. Each man (woman) can match to at most one woman (man) and has a complete strict preference list over all women (men). Given a matching M , a *blocking pair* is a man-woman pair such that both strictly improve by matching to each other. A matching without blocking pair is a *stable matching*.

A central assumption in stable marriage is that every agent knows all agents it can match to. In reality, however, agents often have limited information about their matching possibilities. For instance, in a large society we would not expect a man to match up with any other woman immediately. Instead, there exist restrictions in terms of knowledge and information that allow some pairs to match up directly, while others would have to get to know each other before being able to start a relationship. We incorporate this aspect by assuming that agents are embedded in a fixed network of *links*. Links represent an enduring knowledge relation that is not primarily under the control of the agents. Depending on the interpretation, links could represent, e.g., family, neighbor, colleague or teammate relations. Each agent strives to build one *matching edge* to a partner. The set of links and edges defines a dynamic information structure based on *triadic closure*, a standard idea in social network theory: If two agents have a common friend, they are likely to meet and learn about each other. Translated into our model this implies that each agent can match only to partners in its 2-hop neighborhood of the network of matching edges and links. Then, a *local blocking pair* is a blocking pair of agents that are at hop distance at most 2 in the network. Consequently, a *locally stable matching* is a matching without local blocking pairs. Local blocking pairs are a subset of blocking pairs. In turn, every stable matching is a locally stable matching, because it allows no (local or global) blocking pairs. Thus, one might be tempted to think that locally stable matchings are easier to find and/or reach using distributed dynamics than ordinary stable matchings. In contrast, we show in this paper that locally stable matchings have a rich structure and can behave quite differently than ordinary stable matchings. Our study of locally stable matching with arbitrary strict preferences significantly extends recent work on the special case of correlated or weighted matching [11], in which preferences are correlated via matching edge benefits.

Contribution We concentrate on the important case of two-sided markets, in which a (locally) stable matching is always guaranteed to exist. Our primary interest is to characterize convergence properties of iterative round-based dynamics with distributed control, in which in each round a local blocking pair is resolved. We focus on the REACHABILITY problem: Given an instance and an initial matching, is there a sequence of local blocking pair resolutions leading to a locally stable matching? In Section 3 we see that there are cases, in which a locally stable matching might never be reached. This is in strong contrast to the case of weighted matching, in which it is easy to show convergence of every sequence of local blocking pair resolutions with a potential function. In fact, it

is NP-hard to decide REACHABILITY, even if the network exists only among one partition of agents. Moreover, there exist games and initial matchings such that *every* sequence of local blocking pairs terminating in a locally stable matching is exponentially long. Hence, REACHABILITY might even be outside NP. If we need to decide REACHABILITY for a given initial matching *and a specific locally stable matching to be reached*, the problem is even NP-hard for correlated matching.

Our NP-hardness results hold even if the network exists only among one partition. In Section 4, we concentrate on a more general class of games in which links exist in one partition and between partitions (i.e., one partition has no internal links). This is a natural assumption when considering objects that do not generate knowledge about each other, e.g., when matching resources to networked nodes or users, where initially resources are only known to a subset of users. Here we characterize the impact of memory on distributed dynamics. For *recency memory*, each agent remembers in every round the *most recent partner* that is different from the current one. With recency memory, REACHABILITY is always true, and for every initial matching there exists a sequence of polynomially many local or remembered blocking pairs leading to a locally stable matching. In contrast, for *quality memory* where all agents remember their *best partner* REACHABILITY stays NP-hard. This formally supports the intuition that recency memory is more powerful than quality memory, as the latter yields agents that are “hung up” on preferred but unavailable partners. This provides a novel distinction between recency and quality memory.

Our positive results for recency memory in Section 4 imply that if we pick local blocking pairs uniformly at random in each step, we achieve convergence with probability 1. The proof relies only on the memory of one partition. In contrast, if only the other partition has memory, we obtain NP-hardness of REACHABILITY. Convergence with probability 1 can also be guaranteed for *random memory* if in each round each agent remembers one of his previous matches chosen uniformly at random. The latter result holds even when links exist among or between both partitions. However, using known results on stable marriage with full information [1], convergence time can be exponential with high probability, independently of any memory.

In contrast to ordinary stable matchings, two locally stable matchings can have very different size. This motivates our search for maximum locally stable matchings in Section 5. While a simple 2-approximation algorithm exists, we can show a non-approximability result of $1.5 - \varepsilon$ under the unique games conjecture. For spatial reasons most of the proofs are omitted but can be found in the full version.

Related Work Locally stable matchings were introduced by Arcaute and Vassilvitskii [2] in a two-sided job-market model, in which links exist only among one partition. The paper uses strong uniformity assumptions on the preferences and addresses the lattice structure for stable matchings and a local Gale-Shapley algorithm. More recently, we studied locally stable matching with correlated preferences in the roommates problem, where arbitrary pairs of agents can be matched [11]. Using a potential function argument, REACHABILITY is always

true and convergence guaranteed. Moreover, for every initial matching there is a polynomial sequence of local blocking pairs that leads to a locally stable matching. The expected convergence time of random dynamics, however, is exponential. If we restrict to resolution of pairs with maximum benefit, then for random memory the expected convergence time becomes polynomial, but for recency or quality memory convergence time remains exponential, even if the memory is of polynomial size.

For an introduction to stable marriage and some of its variants we refer the reader to several books in the area [10, 19]. There is a significant literature on dynamics, especially in economics, which is too broad to survey here. These works usually do not address issues like computational complexity or worst-case bounds. We focus on a subset of prominent analytical works related to our scenario. For the stable marriage problem, it is known that better-response dynamics, in which agents sequentially deviate to blocking pairs, can cycle [16]. On the other hand, REACHABILITY is always true, and for every initial matching there exists a sequence of polynomially many steps to a stable matching [20]. If blocking pairs are chosen uniformly at random at each step, convergence time is exponential [1] in the worst case.

In the roommates problem, in which every pair of agents can be matched, stable matchings can be absent. Deciding existence and computing stable matchings if they exist can be done in polynomial time [14]. In addition, if a stable matching exists, then REACHABILITY is always true [8]. A similar statement can be made even more generally for relaxed concepts like P -stable matchings that always exist [12]. Ergodic sets of the underlying Markov chain have been studied [13] and related to random dynamics [15]. In addition, for computing (variants of) stable matchings via iterative entry dynamics see [4–6].

The problem of computing a maximum locally stable matchings has recently been considered in [7]. In addition to characterizations for special cases, a NP-hardness result is shown and non-approximability of $(21/19 - \varepsilon)$ unless $P = NP$. Computing maximum stable matchings with ties and incomplete lists has generated a significant amount of research interest over the past decade. The currently best results are a 1.5-approximation algorithm [17] and $(4/3 - \varepsilon)$ -hardness under the unique games conjecture [21].

2 Preliminaries

A *network matching game* (or *network game*) consists of a (social) *network* $N = (V, L)$, where V is a set of vertices representing *agents* and $L \subseteq \{\{u, v\} | u, v \in V, u \neq v\}$ is a set of fixed *links*. A set $E \subseteq \{\{u, v\} | u, v \in V, u \neq v\}$ defines the *potential matching edges*. A *state* is a matching $M \subseteq E$ such that for each $v \in V$ we have $|\{e \mid e \in M, v \in e\}| \leq 1$. An edge $e = \{u, v\} \in M$ provides utilities $b_u(e), b_v(e) > 0$ for u and v , respectively. If for every $e \in E$ we have $b_u(e) = b_v(e) = b(e) > 0$, we speak of *correlated preferences* or a *correlated network game*. Otherwise, we will assume that each agent has a total order \succ over its possible matching partners, and for every agent the utility of matching edges is given

according to this ranking. Throughout the paper we focus on the *two-sided* or *bipartite* case, which is often referred to as the *stable marriage problem*, where V is divided into two disjoint sets U and W such that $E \subseteq \{\{u, w\} \mid u \in U, w \in W\}$. Note that this does not imply that N is bipartite. If further the vertices of U are isolated in N , we speak of a *job-market game* for consistency with [2, 11].

To describe stability in network matching games, we assume agents u and v are *accessible* in state M if they have a distance of at most 2 in the graph $G = (V, L \cup M)$. A state M has a *local blocking pair* $e = \{u, v\} \in E$ if u and v are accessible and are each either unmatched in M or matched through an edge e' such that e' serves a strictly smaller utility than e . Thus, in a local blocking pair both agents can strictly increase their utility by generating e (and possibly dismissing some other edge thereby). A state M that has no local blocking pair is a *locally stable matching*.

Most of our analysis concerns iterative round-based dynamics, where in each round we pick one local blocking pair, add it to M , and remove all edges that conflict with this new edge. We call one such step a *local improvement step*. With *random dynamics* we refer to the process when in each step the local blocking pair is chosen uniformly at random from the ones available. A local blocking pair $\{u, v\}$ that is resolved and not a link as well must be connected by some distance-2 path (u, w, v) in M before the step. This path can consist of two links, or of exactly one link and one matching edge. In the latter case, let w.l.o.g. $\{u, w\}$ be the matching edge. As u can have only one matching edge, the local improvement step will delete $\{u, w\}$ to create $\{u, v\}$. For simplicity, we will refer to this fact as "an edge moving from $\{u, w\}$ to $\{u, v\}$ " or " u 's edge moving from w to v ".

In subsequent sections we will assume that agents have memory that allows to "remember" one matching partner from a former round. In this case, a pair $\{u, v\}$ of agents becomes accessible not only by a distance-2 path in G , but also when u appears in the memory of v . Hence, in this case a local blocking pair can be based solely on access through memory. For *random memory*, we assume that in every round each agent remembers a previous matching partner chosen uniformly at random. For *recency memory*, each agent remembers the last matching partner that is different from the current partner. For *quality memory*, each agent remembers the best previous matching partner.

3 Complexity of REACHABILITY and Duration

Complexity In contrast to ordinary stable marriage, there exist small examples that show REACHABILITY is not always true for locally stable matchings (see [11, Thm. 5] or our circling gadget in the full version). Here we consider complexity of this problem and show NP-hardness results when agents have strict preferences. This is in contrast to correlated network games, where REACHABILITY is true and for every initial matching there is always a sequence to a locally stable matching of polynomial length [11]. Still, we here show that given a particular matching to reach, deciding REACHABILITY becomes also NP-hard, even for correlated job-

market games. We present the proof of the latter result in detail, as it provides the basic idea for the omitted NP-hardness proofs as well.


Theorem 1. *It is NP-hard to decide REACHABILITY from the initial matching $M = \emptyset$ to a given locally stable matching in a correlated network game.*

Proof. We use a reduction from 3SAT . Given a 3SAT formula with k variables x_1, \dots, x_k and l clauses C_1, \dots, C_l , where clause C_j contains the literals $l1_j, l2_j$ and $l3_j$, we have

$$U = \{u_{x_i} | i = 1 \dots k\} \cup \{u_{C_j} | j = 1 \dots l\} \cup \{b_h | h = 1 \dots k + l - 1\},$$

$$W = \{v_{x_i}, x_i, \bar{x}_i | i = 1 \dots k\} \cup \{v_{C_j} | j = 1 \dots l\} \cup \{a, a_1\}.$$

For the social links see the picture shown below.



We do not restrict the set of matching edges, but assume that every edge not appearing in the list below has benefit $\epsilon \ll 1$ (resulting in them being irrelevant for the dynamics). The other benefits are given as follows.

$u \in U$	$w \in W$	$b(\{u, w\})$	
u_{C_j}	a	j	$j = 1, \dots, l$
u_{x_i}	a	$i + l$	$i = 1, \dots, k$
b_h	a	$h + \frac{1}{2}$	$h = 1, \dots, k + l - 1$
u_{C_j}	$l1_j/l2_j/l3_j$	$k + l + 1$	$j = 1, \dots, l$
u_{x_i}	x_i/\bar{x}_i	$k + l + 1$	$i = 1, \dots, k$
u_{C_j}	v_{x_i}	$k + l + 1 + i$	$i = 1, \dots, k, j = 1, \dots, l$
u_{x_i}	$v_{x_{i'}}$	$k + l + 1 + i'$	$i = 1, \dots, k, i' = 1, \dots, i$
u_{C_j}	$v_{C_{j'}}$	$2k + l + 1 + j'$	$j = 1, \dots, l, j' = 1, \dots, j$

Our goal is to reach $M^* = \{\{u_s, v_s\} | s \in \{x_1, \dots, x_k\} \cup \{C_1, \dots, C_l\}\}$.

First, note that additional matching edges can only be introduced at $\{u_{C_1}, a\}$. Furthermore, once a vertex u_y , $y \in \{x_1, \dots, x_k\} \cup \{C_1, \dots, C_l\}$, is matched to a vertex other than a , it blocks the introduction of any edge for a vertex lying behind u_y on the path from u_{C_1} to u_{x_k} . Also, the vertices b_h prevent that an edge

is moved on from one u -vertex to another after it has left a . Thus, at the time when an edge to a clause u -vertex is created that still exists in the final matching (but is connected to some v_{C_j} then), the edges for all variable u -vertices must have been created already.

Assume that the 3SAT formula is satisfiable. Then we first create a matching edge at $\{u_{C_j}, a\}$, move it over the u - and b -vertices to u_{x_k} , and then move it into the branching to the one of x_k or \bar{x}_k that negates its value in the satisfying assignment. Similarly, one after the other (in descending order), we create a matching edge at a for each of the variable u -vertices and move it into the branching to the variable vertex that negates its value in the satisfying assignment. As every clause is fulfilled, at least one of the three vertices that yield an improvement for the clause u -vertex from a is not blocked by a matching edge to a variable u -vertex. Then, the edges to clause u -vertices can bypass the existing edges (again, one after the other in descending order) and reach their positions in M^* . After that, the variable-edges can leave the branching and move to their final position in the same order as before.

Now assume that we can reach M^* from \emptyset . We note that the edges to clause u -vertices have to overtake the edges to variable u -vertices somewhere on the way to reach their final position. The only place to do so is in the branching leading over the x_i and \bar{x}_i . Thus all variable-edges have to wait at some x_i or \bar{x}_i until the clause-edges have passed. But from a , vertex u_{x_i} is only willing to switch to x_i or \bar{x}_i . Thus, every vertex blocks out a different variable (either in its true or in its false value). Similarly, a vertex u_{C_j} will only move further from a if it can reach one of its literals. Hence, if all clauses can bypass the variables, then for every clause there was one of its literals left open for passage. Thus, if we set each variable to the value that yields the passage for clause-edges in the branching, we obtain a satisfying assignment. \square

Corollary 1. *It is NP-hard to decide REACHABILITY to a given locally stable matching in a correlated job-market game.*

Theorem 2. *It is NP-hard to decide REACHABILITY from the initial matching $M = \emptyset$ to an arbitrary locally stable matching in a bipartite network game.*

Length of Sequences We now consider the number of improvement steps required to reach locally stable matchings. In general, there is a network game and an initial matching such that we need an exponential number of steps before reaching any locally stable matching. This is again in contrast to the correlated case, where every reachable stable matching can be reached in a polynomial number of steps. We present the latter result and defer the much more technical proof of the general lower bound to the full version.

Theorem 3. *For every network game with correlated preferences, every locally stable matching $M^* \in E$ and initial matching $M_0 \in E$ such that M^* can be reached from M_0 through local improvement steps, there exists a sequence of at most $O(|E|^3)$ local improvement steps leading from M_0 to M^* .*

Proof. Consider an arbitrary sequence between M_0 and M^* . We will explore which steps in the sequence are necessary and which parts can be omitted. We rank all edges by their benefit (allowing multiple edges to have the same rank) such that $r(e) > r(e')$ iff $b(e) > b(e')$ and set $r_{max} = \max\{r(e) | e \in E\}$. Recall from Section 2 that we can account edges in the way that every edge e has at most one direct predecessor e' in the sequence, which was necessary to build e . Because e was a local blocking pair, we know $r(e') < r(e)$. Thus, every edge e has at most r_{max} predecessors. Our proof is based on two crucial observations:

- (1) An edge can only be deleted by a stronger edge, that is, every chain of one edge deleting the next is limited in length by r_{max} .
- (2) If an edge is created, then possibly moved, and finally deleted without deleting an edge on its way, this edge would not have to be introduced in the first place.

Suppose our initial matching is the empty matching, then every edge in the locally stable matching has to be created and by (repeated application of) (2) we only need to create and move edges that are needed for the final matching. Thus we have $|M^*|$ edges, which each made at most r_{max} steps.

Now if we start with an arbitrary matching, the sequence might be forced to delete some edges that cannot be used for the final matching. Each of these edges generates a chain of edges deleting each other throughout the sequence, but (1) tells us that this chain is limited as well as the number of steps each of these edges has to make. The only remaining issue is what happens to edges "accidentally" deleted during this procedure. Again, we can use (2) to argue that there is no reason to rebuild such an edge just to delete it again. Thus, such deletions can happen only once for every edge we had in M_0 (not necessarily on the position it had in M_0). It does not do any harm if it happens to an edge of one of the deletion-chains, as it would just end as desired. For the edges remaining in $|M^*|$ the same bounds holds as before. Thus, we have an overall bound of $|M_0| \cdot r_{max} \cdot r_{max} + |M^*| \cdot r_{max} \in O(|E|^3)$ steps, where the first term results from the deletion chains and the second one from the edges surviving in the final matching. \square

Theorem 4. *There is a network game with strict preferences such that a locally stable matching can be reached by a sequence of local improvement steps from the initial matching $M = \emptyset$, but every such sequence has length $2^{\Omega(|V|)}$.*

4 Memory

Given the impossibility results in the last section, we now focus on the impact of memory. As a direct initial result, no memory can yield reachability of a *given* locally stable matching, even in a correlated job-market game.

Corollary 2. *It is NP-hard to decide REACHABILITY to a given locally stable matching in a correlated job-market game with any kind of memory.*

Let us instead concentrate on the impact of memory on reaching an arbitrary locally stable matching. For our treatment we will focus on the case in which the network links $L \subseteq (W \times W) \cup (U \times W)$. We assume that in every step, every agent remembers one previous matching partner.

Quality Memory With quality memory, each agent remembers the best matching partner he ever had before. While this seems quite a natural choice and appears like a smart strategy, it can be easily fooled by starting with a much-liked partner, who soon after matches with someone more preferred and never becomes available again. This way the memory becomes useless which leaves us with the same dynamics as before.

Proposition 1. *There is a network game with strict preferences, links $L \subseteq (W \times W) \cup (U \times W)$, quality memory and initial matching $M = \emptyset$ such that no locally stable matching can be reached with local improvement steps from M . This even holds if every agent remembers the best k previous matches.*

Theorem 5. *It is NP-hard to decide REACHABILITY to an arbitrary locally stable matching in a network game with quality memory.*

Recency Memory With recency memory, each agent remembers the last partner he has been matched to. This is again quite a very natural choice as it expresses the human character of remembering the latest events best. Interestingly, here we actually can ensure that a locally stable matching can be reached.

Theorem 6. *For every network game with strict preferences, links $L \subseteq (U \times W) \cup (W \times W)$, recency memory and every initial matching, there is a sequence of $O(|U|^2|W|^2)$ many local improvement steps to a locally stable matching.*

Proof. Our basic approach is to construct the sequence in two phases similarly as in [1]. In the first phase, we let the matched vertices from U improve, but ignore the unmatched ones. In the second phase, we make sure that vertices from W have improved after every round.

Preparation phase: As long as there is at least one $u \in U$ with u matched and u part of a blocking pair, allow u to switch to the better partner.

The preparation phase terminates after at most $|U| \cdot |W|$ steps, as in every round one matched $u \in U$ strictly improves in terms of preference. This can happen at most $|W|$ times for each matched u . In addition, the number of matched vertices from U only decreases.

Memory phase: As long as there is a $u \in U$ with u part of a blocking pair, pick u and execute a sequence of local improvement steps involving u until u is not part of any blocking pair anymore. For every edge $e = \{u', w\}$ with $u' \neq u$ that was deleted during the sequence, recreate e from the memory of u' .

We claim that if we start the memory phase after the preparation phase, at the end of every round we have the following invariants: The vertices from W that have been matched before are still matched, they do not have a worse

partner than before, and at least one of them is matched strictly better than before. Also, only unmatched vertices from U are involved in local blocking pairs.

Obviously, at the end of the preparation phase the only U -vertices in local blocking pairs are unmatched, i.e., initially only unmatched U -vertices are part of local blocking pairs. Let u be the vertex chosen in the following round of the memory phase. At first we consider the outcome for $w \in W$. If w is the vertex matched to u in the end, then w clearly has improved. Otherwise w gets matched to its former partner (if it had one) through memory and thus has the same utility as before. In particular, every w that represents an improvement to some u' but was blocked by a higher ranked vertex still remains blocked. Together with the fact that u plays local improvement steps until it is not part of a local blocking pair anymore, this guarantees that all matched U -vertices cannot improve at the end of the round. As one W -vertex improves in every round, we have at most $|U| \cdot |W|$ rounds in the memory phase, where every round consists of at most $|W|$ steps by u and at most $|U| - 1$ edges reproduced from memory. \square

The existence of sequences to (locally) stable matchings also implies that random dynamics converge in the long run with probability 1 [8, 12, 20]. In general, we cannot expect fast convergence here, as there are instances where random dynamics yield an exponential sequence with high probability even if all information is given – e.g., reinterpret the instance from [1] with $L = U \times W$, then every agent knows every possible partner and memory has no effect.

Observe that the previous proof relies only on the recency memory of partition U . Hence, the existence of short sequences holds even if only agents from U have memory. In contrast, if only agents from W have recency memory, the previous NP-hardness constructions can be extended.

Theorem 7. *It is NP-hard to decide REACHABILITY to an arbitrary locally stable matching in a network game with links $L \subseteq (U \times W) \cup (W \times W)$ and recency memory only for agents in W .*

Random Memory Finally, with random memory, each agent remembers a partner chosen uniformly at random in each step. We consider random memory and reaching a locally stable matching from every starting state even in general network games. While we cannot expect fast convergence, we can show that random memory helps with reachability:

Theorem 8. *For every network game with random memory, random dynamics converge to a locally stable matching with probability 1.*

5 Maximum Locally Stable Matchings

As the size of locally stable matchings can vary significantly – up to the point where the empty matching as well as a matching that includes every vertex is locally stable – it is desirable to target locally stable matchings of maximal size.

We address the computational complexity of finding maximum locally stable matchings by relating it to the independent set problem.

Theorem 9. *For every graph $G = (V, E)$ there is a job-market game that admits a maximum locally stable matching of size $|V| + k$ if and only if G holds a maximum independent set of size k .*

Proof. Given a graph $G = (V, E)$, $|V| = n$, we construct the job-market game with network $N = (V' = U \cup W, L)$. For every $v \in V$ we have $u_{v,1}, u_{v,2} \in U$ and $w_{v,1}, w_{v,2} \in W$. We have the links $\{w_{v,1}, w_{v',2}\}$ and $\{w_{v',2}, w_{v,2}\}$ if $v' \in N(v)$. We allow matching edges $\{u_{v,1}, w_{v,1}\}$, $\{u_{v,1}, w_{v',2}\}$ for $v' \in N(v)$, $\{u_{v,1}, w_{v,2}\}$ and $\{u_{v,2}, w_{v,2}\}$. Each $u_{v,1}$ prefers $w_{v,2}$ to every $w_{v',2}$, $v' \in N(v)$, and every $w_{v',2}$ to $w_{v,1}$. The preferences between the different neighbors can be chosen arbitrarily. Each $w_{v,2}$ prefers $u_{v,1}$ to every $u_{v',1}$, $v' \in N(v)$, and every $u_{v',2}$ to $u_{v,2}$. Again the neighbors can be ordered arbitrarily. Vertices $w_{v,1}$ and $u_{v,2}$ have only one possible matching partner.

We claim that G has a maximum independent set of size k iff N has a locally stable matching of size $n + k$.

Let S be a maximum independent set in G . Then $M = \{\{u_{v,1}, w_{v,2}\} \mid v \in V \setminus S\} \cup \{\{u_{v,1}, w_{v,1}\}, \{u_{v,2}, w_{v,2}\} \mid v \in S\}$ is a locally stable matching as the edges $\{u_{v,1}, w_{v,2}\}$ are always stable. For the other vertices the independent set property tells us that for $v \in S$ all vertices $v' \in N(S)$ generate stable edges $\{u_{v',1}, w_{v',2}\}$ that keep $u_{v,1}$ from switching to $w_{v',2}$. Thus $\{u_{v,1}, w_{v,1}\}$ is stable and $w_{v,2}$ cannot see $u_{v,1}$ which stabilizes $\{u_{v,2}, w_{v,2}\}$.

Now let M be a maximum locally stable matching for the job-market game. Further we chose M such that every $u_{v,1}$ is matched, which is possible as replacing a matching partner of $w_{v,2}$ by (the unmatched) $u_{v,1}$ will not generate instabilities or lower the size of M . We note that no $u_{v,1}$ is matched to some $w_{v',2}$ with $v \neq v'$ as from there $u_{v,1}$ and $w_{v,2}$ can see each other and thus constitute a blocking pair. Then, for $S = \{v \mid u_{v,2} \in M\}$, $|S| = |M| - n$ and S is an independent set, as every $u_{v,2}$ can only be matched to its vertex $w_{v,2}$, which means that $u_{v,1}$ must be matched to $w_{v,1}$. But this edge is only stable if every $w_{v',2}$, $v' \in N(v)$, is blocked by $u_{v',1}$. Hence for every $v \in S$ $N(v) \cap S = \emptyset$. \square

This result allows us to transfer hardness of approximation results for independent set to locally stable matching.

Corollary 3. *Finding a maximum locally stable matching is NP-complete. Under the unique games conjecture the problem cannot be approximated within $1.5 - \varepsilon$, for any constant ε .*

In fact, our reduction applies in the setting of the job-market game, where one side has no network at all. This shows that even under quite strong restrictions the hardness of approximation holds. In contrast, it is easy to obtain a 2-approximation in every network game that admits a globally stable matching.

Proposition 2. *If a (globally) stable matching exists, every such stable matching is a 2-approximation for the maximum locally stable matching.*

References

1. H. Ackermann, P. Goldberg, V. Mirrokni, H. Röglin, and B. Vöcking. Uncoordinated two-sided matching markets. *SIAM J. Comput.*, 40(1):92–106, 2011.
2. E. Arcaute and S. Vassilvitskii. Social networks and stable matchings in the job market. In *Proc. 5th Intl. Workshop Internet & Network Economics (WINE)*, pages 220–231, 2009.
3. P. Austrin, S. Khot, and M. Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. *Theory of Computing*, 7(1):27–43, 2011.
4. P. Biró, K. Cechlárová, and T. Fleiner. The dynamics of stable matchings and half-matchings for the stable marriage and roommates problems. *Int. J. Game Theory*, 36(3–4):333–352, 2008.
5. Y. Blum, A. Roth, and U. Rothblum. Vacancy chains and equilibration in senior-level labor markets. *J. Econom. Theory*, 76:362–411, 1997.
6. Y. Blum and U. Rothblum. “Timing is everything” and martial bliss. *J. Econom. Theory*, 103:429–442, 2002.
7. C. Cheng and E. McDermid. Maximum locally stable matchings. In *Proc. 2nd Intl. Workshop Matching under Preferences (MATCH-UP)*, pages 51–62, 2012.
8. E. Diamantoudi, E. Miyagawa, and L. Xue. Random paths to stability in the roommates problem. *Games Econom. Behav.*, 48(1):18–28, 2004.
9. M. Goemans, L. Li, V. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. *IEEE J. Sel. Area Comm.*, 24(5):1020–1033, 2006.
10. D. Gusfield and R. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.
11. M. Hoefer. Local matching dynamics in social networks. *Inf. Comput.*, 222:20–35, 2013.
12. E. Inarra, C. Larrea, and E. Moris. Random paths to P -stability in the roommates problem. *Int. J. Game Theory*, 36(3–4):461–471, 2008.
13. E. Inarra, C. Larrea, and E. Moris. The stability of the roommate problem revisited. Core Discussion Paper 2010/7, 2010.
14. R. Irving. An efficient algorithm for the “stable roommates” problem. *J. Algorithms*, 6(4):577–595, 1985.
15. B. Klaus, F. Klijn, and M. Walzl. Stochastic stability for roommate markets. *J. Econom. Theory*, 145:2218–2240, 2010.
16. D. Knuth. *Marriages stables et leurs relations avec d’autres problemes combinatoires*. Les Presses de l’Université de Montréal, 1976.
17. E. McDermid. A $3/2$ -approximation algorithm for general stable marriage. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 689–700, 2009.
18. A. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *J. Econom. Theory*, 125(2):151–188, 2005.
19. A. Roth and M. O. Sotomayor. *Two-sided Matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
20. A. Roth and J. V. Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, 1990.
21. H. Yanagisawa. *Approximation algorithms for stable marriage problems*. PhD thesis, Kyoto University, Graduate School of Informatics, 2007.