

Article

Maximum Locally Stable Matchings

Christine T. Cheng^{1,*} and Eric McDermid²

¹ Department of Computer Science, University of Wisconsin-Milwaukee, 3200 N. Cramer Street, Milwaukee, WI 53211, USA

² 21CT, 6011 West Courtyard Drive, Austin, TX 78730, USA; E-Mail: emcdermid@21ct.com

* Author to whom correspondence should be addressed; E-Mail: ccheng@uwm.edu; Tel.: +1-414-229-5170; Fax: +1-414-229-2769.

Received: 4 January 2013; in revised form: 2 June 2013 / Accepted: 4 June 2013 /

Published: 24 June 2013

Abstract: Motivated by the observation that most companies are more likely to consider job applicants referred by their employees than those who applied on their own, Arcaute and Vassilvitskii modeled a job market that integrates social networks into stable matchings in an interesting way. We call their model HR+SN because an instance of their model is an ordered pair (I, G) where I is a typical instance of the Hospital/Residents problem (HR) and G is a graph that describes the social network (SN) of the residents in I . A matching μ of hospitals and residents has a *local blocking pair* (h, r) if (h, r) is a blocking pair of μ , and there is a resident r' such that r' is simultaneously an employee of h in the matching and a neighbor of r in G . Such a pair is likely to compromise the matching because the participants have access to each other through r' : r can give her resume to r' who can then forward it to h . A *locally stable matching* is a matching with no local blocking pairs. The cardinality of the locally stable matchings of I can vary. This paper presents a variety of results on computing a locally stable matching with maximum cardinality.

Keywords: stable matchings; hospital/residents problem; social networks

1. Introduction

Motivated by the observation that most companies are more likely to consider job applicants suggested by their employees than those who apply on their own, Arcaute and Vassilvitskii [1] modeled a job market that integrates social networks into stable matchings. Formally, an instance of their model

consists of a set of firms F , a set of workers W , and a social network graph G of the workers. Each member of $F \cup W$ has a preference list that ranks members of the opposite group that it or she finds acceptable in some linear order. Each firm f has a capacity q_f , the maximum number of workers it can employ. A firm-worker pair is *acceptable* if they appear in each other's preference lists. A (many-to-one) *matching* μ of F and W is a set of acceptable firm-worker pairs where each firm f is part of at most q_f pairs and each worker is part of at most one pair. It has a *blocking pair* (f, w) if (i) (f, w) is an acceptable pair; (ii) f has an opening or f prefers w to its worst employee under μ ; and (iii) w is unemployed or w prefers f to her employer under μ . This blocking pair is *local* if additionally f and w have access to each other—i.e., f has an employee that is also a neighbor of w in G . In two-sided matching theory, a common goal is to find *stable matchings*, which are matchings with no blocking pairs. In this model, the matchings of interest are *locally stable matchings*, which are matchings with no local blocking pairs.

We emphasize that a locally stable matching can contain blocking pairs, but these blocking pairs are unlikely to compromise the matching. This may seem odd—if f still has an opening or prefers w over its worst employee, why f cannot just make a job offer to w ? If w is unemployed or prefers f over her current employer, why she cannot just apply to f ? In the job market context, we can think of the preference lists of f and w as being constructed in an online fashion. In particular, w can be included in f 's list only *after* f has seen w 's resume. However, for f to consider w , some employee of f who is also a friend of w must forward w 's resume to f . Similarly, f can be included in w 's list only *after* w has gained some reliable information about f . For this to happen, w must have a friend who works at f . This friend of w that is also an employee of f is a *point of contact* between f and w . The main assumption in this model is that a blocking pair of a matching cannot affect the matching if the firm-worker pair has no points of contact. Note, however, that although a firm-worker pair (f, w) may exist such that w has no social ties with any of f 's other acceptable partners, (f, w) can still be involved in a matching.

The first part of Arcaute and Vassilvitskii's paper [1] explores the combinatorial differences between stable matchings and locally stable matchings. Among others, they show that there are instances whose set of locally stable matchings do not form a distributive lattice under the standard ordering relation used for stable matchings. There are also instances whose locally stable matchings vastly outnumber its stable matchings. The second part of their paper examines the evolution of the job market. They consider a decentralized version of Gale and Shapley's algorithm and show that for a specific case the algorithm converges to a locally stable matching under weak stochastic conditions. They then go on to analyze the goodness of the resulting locally stable matching. The recent work of Hoefer [2] expands on the latter line of inquiry significantly.

Our Contribution. In this paper, we continue the study of locally stable matchings, focusing on those with maximum cardinality. We call them the *maximum locally stable matchings*. In our opinion, not only are the locally stable matchings inherently interesting but they are also an intriguing alternative to stable matchings. In some applications, requiring a matching to be stable can be too strong a requirement. It can also unnecessarily limit the size of the matching. This has led researchers to suggest other kinds of matchings that still take participants' preferences into consideration. They include popular matchings [3] and its many variants (e.g., [4–6], etc.), rank maximal matchings [7], and “almost stable” maximum matchings—which are maximum matchings with few blocking pairs [8]. In the job market context,

locally stable matchings may not only be larger than stable matchings, they may be just as robust since participants are unlikely to leave their assignments. Here are our main contributions:

- First, we present families of instances where the problem of finding a maximum locally stable matchings is computationally easy. For one family of instances, every stable matching of the instance is a maximum locally stable matching. This family includes the case when G , the social network of the workers, is a complete graph. For the other family of instances, every maximum matching of the firms and workers is a maximum locally stable matching. This family includes the case when G is an empty graph.
- Next, we show that when \bar{G} , the complement of G , has a maximum matching of size r , the size of a maximum locally stable matching of the instance is at most r more than the size of a stable matching of the instance. Thus, when G is almost a complete graph, a stable matching of the instance is a good approximation to its maximum locally stable matching. On the other hand, we show that when G has a constant number of edges—*i.e.*, G is almost an empty graph—finding a maximum locally stable matching can still be done in polynomial time.
- Finally, in spite of the results above, we show that finding a maximum locally stable matching is computationally hard in general. In particular, we prove that finding a locally stable matching of a certain size is NP-complete and that approximating the size of a maximum locally stable matching within $21/19 - \delta$ is NP-hard. Recently, Hoefer and Wagner [9] have shown that this problem cannot be approximated within $1.5 - \epsilon$ under the Unique Games Conjecture. Hence, our result differs from theirs as we require only the weaker assumption that $P \neq NP$.

The rest of the paper is organized as follows. In Section 2, we state facts and preliminary results. We present the first two results in Section 3, and the last result in Section 4. We conclude in Section 5.

2. Preliminaries

In the stable matchings literature, the problem of finding a stable matching in the Arcaute–Vassilvitskii model sans the social network is often referred to as the Hospital/Residents problem (HR). The firms correspond to the hospitals while the workers correspond to the residents. In their seminal paper on stable matchings [10], Gale and Shapley presented an algorithm that finds a stable matching for every HR instance I . It can be implemented in $O(|I|)$ time where $|I|$ is the number of acceptable pairs in I . In general, I can have many stable matchings. Nonetheless, Gale and Sotomayor [11] showed that every stable matching of I has the same size and matches exactly the same set of residents. Throughout this paper, we shall assume that every HR instance we can consider has the property that a resident r is in a hospital h 's preference list if and only if h is also in r 's preference list.

An example. In the following instance, let the hospitals be h_1, h_2, h_3 whose capacities are 2, 2, 4 respectively. Let the set of residents be $r_1, r_2, r_3, r_4, r_5, r_6, r_7, r_8$. Here are their preference lists:

$h_1:$	r_1	r_2	r_5	r_6	$r_1:$	h_1	h_3	$r_5:$	h_1
$h_2:$	r_3	r_4	r_7	r_8	$r_2:$	h_1	h_3	$r_6:$	h_1
$h_3:$	r_1	r_2	r_3	r_4	$r_3:$	h_2	h_3	$r_7:$	h_2
					$r_4:$	h_2	h_3	$r_8:$	h_2

It is not difficult to see that $\mu = \{(h_1, r_1), (h_1, r_2), (h_2, r_3), (h_2, r_4)\}$ is a stable matching of the instance. It is, however, smaller than $\sigma = \{(h_1, r_5), (h_1, r_6), (h_2, r_7), (h_2, r_8), (h_3, r_1), (h_3, r_2), (h_3, r_3), (h_3, r_4)\}$, which is a maximum matching of the instance.

For HR instance I , let $B[I]$ denote the bipartite graph where the hospitals are the vertices on one side, and the residents on the other side. A pair (h, r) is an edge if and only if they form an acceptable pair. Thus, every matching of I is a subgraph of $B[I]$. Finding a maximum matching of I can be done by solving a maximum flow problem with $B[I]$ as the “base graph”: Create a source s and a directed edge from s to every hospital h , and set its capacity to q_h . Direct all edges $\{h, r\}$ in $B[I]$ from h to r and set its capacity to 1. Finally, create a sink t and a directed edge from every resident r to t , and set its capacity to 1. There is a one-to-one correspondence between the maximum flows of this network and the maximum matchings of $B[I]$.

Let us suppose that all hospitals in I have capacity 1, and we wish to compare two of its matchings μ and σ . In this case, it is useful to consider their symmetric difference $\mu \oplus \sigma$. In $B[I]$, it is made up of what are called *alternating paths and cycles*—i.e., paths and cycles whose edges alternately belong to μ and σ . In cycles and even-length alternating paths, the number of μ -edges and the number of σ -edges are the same; in odd-length alternating paths, the numbers differ by 1. Additionally, when $|\sigma| > |\mu|$, there is at least one odd-length alternating path with one more σ -edge than μ -edge. We shall call such a path a σ -alternating path.

When not all hospitals in I have capacity 1, we can transform I to another instance where this is the case. Here is a standard trick [12]. Denote by $I^{(1)}$ the instance obtained from I by doing the following: for each hospital h_i of I with capacity q_{h_i} , replace h_i by q_{h_i} clones of h_i : $h_{i,1}, h_{i,2}, \dots, h_{i,q_{h_i}}$. Let their capacities be 1, and let their preference lists be exactly the same as that of h_i . Then for each resident r_j that has h_i in her preference list, replace h_i with the linear order $(h_{i,1}, h_{i,2}, \dots, h_{i,q_{h_i}})$. By transforming I to $I^{(1)}$, the many-to-one matchings of I can now be viewed as one-to-one matchings of $I^{(1)}$. Let μ be a matching of I . Create the corresponding matching $\mu^{(1)}$ of $I^{(1)}$ as follows: when h_i is matched to residents $r_{j_1}, r_{j_2}, \dots, r_{j_k}$ in μ and these residents are arranged according to its preference, let $h_{i,1}, h_{i,2}, \dots, h_{i,k}$ be matched to $r_{j_1}, r_{j_2}, \dots, r_{j_k}$ respectively in $\mu^{(1)}$. Notice that μ and $\mu^{(1)}$ have the same size. Moreover, it is easy to verify that this mapping is a bijection from the set of stable matchings of I to the set of stable matchings of $I^{(1)}$. Now, suppose we want to compare μ with another matching σ of I . The task becomes equivalent to comparing $\mu^{(1)}$ and $\sigma^{(1)}$ in $I^{(1)}$, and the symmetric difference technique described in the previous paragraph can now be applied.

Proposition 1. *In the HR instance I , let μ be a stable matching and σ be a maximum matching of I . Then $|\mu| \leq |\sigma| \leq 2|\mu|$.*

Proof. By definition, $|\mu| \leq |\sigma|$. Now, construct $I^{(1)}$ and the matchings $\mu^{(1)}$ and $\sigma^{(1)}$ from I , μ and σ respectively. Since $|\mu| = |\mu^{(1)}|$ and $|\sigma| = |\sigma^{(1)}|$, $|\mu^{(1)}| \leq |\sigma^{(1)}|$. This means that in $\mu^{(1)} \oplus \sigma^{(1)}$ there is a $\sigma^{(1)}$ -alternating path. However, there cannot be a $\sigma^{(1)}$ -alternating path that simply consists of one edge (h, r) from $\sigma^{(1)}$ because this means that h and r are acceptable to each other in $I^{(1)}$ but are unmatched in $\mu^{(1)}$ —i.e., $\mu^{(1)}$ is not a stable matching of $I^{(1)}$ since (h, r) is a blocking pair. This contradicts the fact that $\mu^{(1)}$ was constructed from μ , a stable matching of I . Thus, in every $\sigma^{(1)}$ -alternating path, the ratio of edges belonging to $\sigma^{(1)}$ to those belonging to $\mu^{(1)}$ is at most 2 : 1. Hence, $|\sigma^{(1)}| \leq 2|\mu^{(1)}|$. It follows that $|\sigma| \leq 2|\mu|$. □

The previous example shows that the bound in Proposition 1 is tight.

2.1. HR+SN and Max-HR+SN

Following the above terminology, we shall call the problem of finding a locally stable matching and a maximum locally stable matching in the Arcaute–Vassilvitskii’s model HR+SN and max-HR+SN respectively, where SN stands for *social network*. We will, however, revert back to the original context and use firms in place of hospitals and workers in place of residents. An instance of HR+SN is an ordered pair (I, G) where I is an HR instance and G is a social network of the workers.

Example continued. In the previous example, suppose G consists of two cliques, one containing r_1, r_2, r_3, r_4 and another containing r_5, r_6, r_7, r_8 . Then σ is a maximum locally stable matching of (I, G) . It has several blocking pairs— $(h_1, r_1), (h_1, r_2), (h_2, r_3), (h_2, r_4)$ —but none of the pairs have a point of contact.

Proposition 2. *If μ is a locally stable matching in the HR+SN instance (I, G) , then $\mu^{(1)}$ is a locally stable matching in the HR+SN instance $(I^{(1)}, G)$.*

Proof. Assume $\mu^{(1)}$ is not a locally stable matching of $(I^{(1)}, G)$ so it has a local blocking pair (f_{ik}, w_j) . Thus, f_{ik} and w_j have a point of contact, say w' , who is also the only employee of f_{ik} in $\mu^{(1)}$. If w_j is unmatched or is employed by a firm that is not a clone of f_i in $\mu^{(1)}$, then (f_i, w_j) is a local blocking pair of μ with w' as a point of contact—a contradiction. If w_j is employed by a clone of f_i in $\mu^{(1)}$, say $f_{ik'}$, then the fact that w_j prefers f_{ik} over $f_{ik'}$ means that $k < k'$. On the other hand, f_{ik} prefers w_j over its only employee w' means that f_i prefers w_j over w' . However, this contradicts the way $\mu^{(1)}$ is constructed because f_i should prefer the worker matched to f_{ik} over the worker matched to $f_{ik'}$. Hence, we have shown that all cases lead to a contradiction. Therefore, $\mu^{(1)}$ is a locally stable matching of $(I^{(1)}, G)$. \square

We note though that the converse of Proposition 2 is not always true as shown by this simple instance $(I^{(1)}, G)$ where

$$\begin{array}{ll} f_{11}: & w_1 \quad w_2 & w_1 : & f_{11} \quad f_{12} \\ f_{12}: & w_1 \quad w_2 & w_2 : & f_{11} \quad f_{12} \end{array}$$

and G consists of the edge w_1w_2 . Consider $\mu^{(1)} = \{(f_{11}, w_1)\}$. It has a blocking pair (f_{12}, w_2) but f_{12} and w_2 have no points of contact because f_{12} has no employee. Hence, $\mu^{(1)}$ is a locally stable matching of $(I^{(1)}, G)$. Now, $I^{(1)}$ corresponds to the instance I where f_1 has a capacity of two, f_1 has the same preference list as f_{11} and f_{12} and both w_1 and w_2 have a preference list with only f_1 in it. Thus, $\mu^{(1)}$ corresponds to the matching $\mu = \{(f_1, w_1)\}$. The latter though is not a locally stable matching of (I, G) because (f_1, w_2) is a local blocking pair with w_1 as a point of contact.

In the next proposition, we provide a bound similar to Proposition 1.

Proposition 3. *In the HR+SN instance (I, G) , let μ be a stable matching and $\hat{\mu}$ be a maximum locally stable matching. Then $|\mu| \leq |\hat{\mu}| \leq 2|\mu|$.*

Proof. Let σ be a maximum matching of I . By definition, $|\mu| \leq |\hat{\mu}| \leq |\sigma|$. According to Proposition 1, $|\sigma| \leq 2|\mu|$. Hence, $|\hat{\mu}| \leq 2|\mu|$. \square

When we appended our running example with the social network consisting of a clique containing r_1, r_2, r_3, r_4 and another clique containing r_5, r_6, r_7, r_8 , σ is a maximum locally stable matching. Its size is twice that of μ . This shows that the bound of Proposition 3 is tight. The next proposition describes the interaction between the preference lists in I and the edges in G .

Proposition 4. *Let (I, G) be an HR instance. Suppose two workers w_1 and w_2 do not have a firm in common in their preference list or, equivalently, there is no firm that has w_1 and w_2 in its preference list. Let $e = \{w_1, w_2\}$. Then $(I, G - e)$ and $(I, G + e)$ have the same set of locally stable matchings as (I, G) .*

Proof. Without loss of generality, assume e is an edge of G . It is easy to verify that when G' is a subgraph of G , every locally stable matching of (I, G) is also a locally stable matching of (I, G') . Thus, to prove the proposition, we simply have to show that every locally stable matching of $(I, G - e)$ is also a locally stable matching of (I, G) .

Suppose μ is a locally stable matching of $(I, G - e)$ but has a local blocking pair (f, w) in (I, G) . Let f and w 's point of contact be w' . Hence, both w and w' have f in their preference lists; that is, $\{w, w'\} \neq \{w_1, w_2\}$. Thus, the edge $\{w, w'\}$ is in $G - e$ so that (f, w) is also a local blocking pair of μ in $(I, G - e)$, a contradiction. It follows that μ has no blocking pairs in (I, G) .

We have shown that μ is a locally stable matching of (I, G) if and only if it is also a locally stable matching of $(I, G - e)$. The same argument holds for the locally stable matchings of (I, G) and $(I, G + e)$ when e is not an edge of G . □

In Section 4, we shall consider max-HR+SN. Given an HR+SN instance (I, G) , let A be an algorithm that outputs a locally stable matching of (I, G) , which we denote as $A(I, G)$. Then A is an $f(N)$ -approximation algorithm of max-HR+SN if for all instances (I, G) of size N , $|\mu^*(I, G)|/|A(I, G)| \leq f(N)$ where $\mu^*(I, G)$ is a maximum locally stable matching of (I, G) . Thus, according to Proposition 3, the Gale–Shapley algorithm is a 2-approximation algorithm of max-HR+SN. The problem max-HR+SN is NP-hard to approximate within $f(N)$ if the existence of an efficient $f(N)$ -approximation algorithm implies $P = NP$.

3. The Easy Cases

In this section, we present a family of HR+SN instances where finding a maximum locally stable matching can be solved in polynomial time.

Theorem 1. *Let (I, G) be an HR+SN instance. Suppose that whenever two workers have a firm in common in their preference lists, the two workers also share an edge in G . Then every stable matching of I is a maximum locally stable matching of (I, G) . Consequently, when G is the complete graph, every stable matching of I is a maximum locally stable matching of (I, G) .*

Proof. For now, assume that all firms in I have capacity 1. Suppose (I, G) has a locally stable matching σ that is larger than the stable matchings of I . Let μ be one of these stable matchings. Then $\sigma \oplus \mu$ has a σ -alternating path of the form $f_1, w_1, f_2, w_2, \dots, f_k, w_k$ such that $(f_i, w_i) \in \sigma$ for $i = 1, \dots, k$ and $(f_{i+1}, w_i) \in \mu$ for $i = 1, \dots, k - 1$. Since f_1 is unmatched in μ , w_1 must prefer f_2 over f_1 ; otherwise, (f_1, w_1) is a blocking pair of μ . Now w_1 and w_2 both have f_2 in their preference lists so they share an

edge in G . It must be the case then that f_2 prefers w_2 over w_1 ; otherwise, (f_2, w_1) is a local blocking pair of σ . Continuing in this fashion, we have that for $i = 1, \dots, k - 1$, w_i prefers f_{i+1} over f_i because μ is a stable matching of I while f_{i+1} prefers w_{i+1} over w_i because w_i and w_{i+1} are adjacent in G and σ is a locally stable matching of (I, G) . Consequently, f_k must prefer w_k over w_{k-1} . However, w_k is unmatched in μ so this implies that (f_k, w_k) is a blocking pair of μ —a contradiction. Hence, σ cannot exist, and μ is a maximum locally stable matching of (I, G) .

So suppose some firms in I have capacity greater than 1. Construct the HR+SN instance $(I^{(1)}, G)$ from (I, G) . Notice that the property “whenever two workers have a firm in common in their preference lists, the two workers also share an edge in G ” is preserved in $(I^{(1)}, G)$. Let σ be a locally stable matching of (I, G) , and let μ be a stable matching of I . Consider their corresponding matchings $\sigma^{(1)}$ and $\mu^{(1)}$. From Proposition 2, $\sigma^{(1)}$ is also a locally stable matching of $(I^{(1)}, G)$. We also know that $\mu^{(1)}$ is a stable matching of $I^{(1)}$. If $|\sigma| > |\mu|$, $|\sigma^{(1)}| > |\mu^{(1)}|$. However, all firms in $I^{(1)}$ have capacity 1, and according to the previous paragraph $\mu^{(1)}$ is a maximum locally stable matching of $(I^{(1)}, G)$. Hence, $|\sigma| \leq |\mu|$, and μ is a maximum locally stable matching of (I, G) . It is easy to see that, as a consequence of Proposition 4, if G is extended to a complete graph, the result still holds. \square

The next theorem provides a bound that is different from the one presented in Proposition 3. It shows that when G is almost a complete graph, a stable matching of I and a maximum locally stable matching of (I, G) will almost have the same size.

Theorem 2. *Let (I, G) be an HR+SN instance. Suppose that the size of the largest matching in \bar{G} , the complement of G , is r . Let $\hat{\mu}$ be a maximum locally stable matching of (I, G) and μ be a stable matching of I . Then $|\hat{\mu}| \leq |\mu| + r$.*

Proof. Once again, let us begin the proof by assuming that all firms in I have capacity 1. First, notice that $|\hat{\mu}| - |\mu|$ is bounded above by the number of $\hat{\mu}$ -alternating paths in $\hat{\mu} \oplus \mu$. Furthermore, the $\hat{\mu}$ -alternating paths are vertex-disjoint. Now, in the proof of Theorem 1, we argued that when $f_1, w_1, f_2, w_2, \dots, f_k, w_k$ forms a $\hat{\mu}$ -alternating path and w_1, w_2, \dots, w_k is a path in G , a contradiction arises. Thus, at least one of the edges $w_1w_2, w_2w_3, \dots, w_{k-1}w_k$ must be missing from G and therefore present in \bar{G} . If $\hat{\mu} \oplus \mu$ has x $\hat{\mu}$ -alternating paths, \bar{G} has at least x pairwise vertex-disjoint edges. Since $x \leq r$, $|\hat{\mu}| - |\mu| \leq r$.

Now suppose some firms in I have capacity greater than 1. Again, construct the HR+SN instance $(I^{(1)}, G)$ from (I, G) . Since $\hat{\mu}$ is a maximum locally stable matching of (I, G) , $\hat{\mu}^{(1)}$ is a locally stable matching of $(I^{(1)}, G)$, but it may not be the largest such matching. Let τ be a maximum locally stable matching of $(I^{(1)}, G)$. Thus, $|\hat{\mu}| - |\mu| = |\hat{\mu}^{(1)}| - |\mu^{(1)}| \leq |\tau| - |\mu^{(1)}|$. From the previous paragraph, the latter is bounded by r . Hence, $|\hat{\mu}| - |\mu| \leq r$. \square

We now consider the opposite case of Theorem 1.

Theorem 3. *Let (I, G) be an HR+SN instance. Suppose that whenever two workers have a firm in common in their preference lists, the two workers do not share an edge in G . Then the matchings of I are exactly the locally stable matchings of (I, G) . Hence, every maximum matching of I is a maximum locally stable matching of (I, G) . Consequently, when G is the empty graph, the matchings of I are exactly the locally stable matchings of (I, G) , and every maximum matching of I is a maximum locally stable matching of (I, G) .*

Proof. Let μ be an arbitrary matching of I . Suppose μ is not stable and contains a blocking pair (f, w) . In order for (f, w) to be a local blocking pair, f and w must have a point of contact w' ; i.e., both w and w' have f in their preference lists, and both are neighbors in G . However, by our assumption on G , this cannot be the case. Hence, all blocking pairs of μ are not local so μ is a locally stable matching of (I, G) . Thus, every matching of I is a locally stable matching of (I, G) . Since every locally stable matching of (I, G) is also a matching of I , it follows that the matchings of I are exactly the locally stable matchings of (I, G) . The rest of the theorem follows. \square

In this next theorem, we consider the case when G is almost an empty graph.

Theorem 4. *Suppose that in the HR+SN instance (I, G) , G has a constant number of edges. Then a maximum locally stable matching of (I, G) can be found in time polynomial in $|I|$.*

Proof. Let $W_1 \subseteq W$ be the smallest set of workers whose induced subgraph G_1 in G contains all the r edges of G . Let $W_2 = W - W_1$, and let G_2 be the subgraph induced by W_2 , which in this case is an empty graph. Thus, $G = G_1 \cup G_2$. Furthermore, every matching μ of I can be expressed as $\mu_1 \cup \mu_2$ where each μ_i is a matching involving the workers in W_i , for $i = 1, 2$. Let I_1 be the HR instance derived from I by restricting the set of workers to W_1 . With some abuse in notation, let $I - \mu_1$ be the HR instance obtained from I by removing the workers matched in μ_1 and decreasing the capacities of the firms according to the number of matches they received in μ_1 . Thus, μ_1 is a matching of I_1 and μ_2 is a matching of $I - \mu_1$. Conversely, if μ_1 is a matching of I_1 and μ_2 is a matching of $I - \mu_1$, putting them together as $\mu = \mu_1 \cup \mu_2$ results in a matching of I .

It is also straightforward to verify that when μ is a locally stable matching of (I, G) , μ_1 and μ_2 are locally stable matchings of (I_1, G_1) and $(I - \mu_1, G_2)$ respectively. Let us now argue the converse. Suppose μ_1 and μ_2 are locally stable matchings of (I_1, G_1) and $(I - \mu_1, G_2)$ respectively but μ has a local blocking pair (f, w) whose point of contact is w_1 . Thus, in μ , f has an opening or prefers w over its worst employee w_2 , and that w is either unmatched or prefers f over her current employer. Furthermore, since G_2 is an empty graph, w and w_1 must both be in W_1 and neighbors in G_1 . If f has an opening in μ or w_2 is in W_2 , then f has an opening after the matching μ_1 so (f, w) is a local blocking pair of μ_1 . On the other hand, if w_2 is in W_1 , then f prefers w to its worst employee w_2 in μ_1 so that (f, w) is again a local blocking pair of μ_1 . All cases lead to a contradiction. Thus, μ must be a locally stable matching of (I, G) .

To find a maximum locally stable matching of (I, G) , we do what is essentially a brute force method. We consider all possible matchings of I_1 . For each such matching μ_1 , we check to see if μ_1 is a locally stable matching of (I_1, G_1) . If it is, we construct $I - \mu_1$ and then find a maximum matching μ_2 of the instance. According to Theorem 3, μ_2 is a maximum locally stable matching of $(I - \mu_1, G_2)$ since G_2 is an empty graph. If $\mu = \mu_1 \cup \mu_2$ is currently the largest locally stable matching of (I, G) we have seen, we store μ ; otherwise, we move on to the next matching of I_1 .

There are at most $2r$ workers in W_1 . Each one is either unmatched or employed by one of the $|F|$ firms. Thus, the number of possible matchings of I_1 is $O((|F| + 1)^{2r})$. Verifying if a matching μ_1 of I_1 is locally stable in (I_1, G_1) , constructing $I - \mu_1$, and finding a maximum matching of the instance can all be done in $O(\text{poly}(|F|, |W|))$ time. Hence, finding a maximum locally stable matching of (I, G) takes $O((|F| + 1)^{2r} \text{poly}(|F|, |W|))$ time, which is polynomial in $|F|$ and $|W|$ when r is a constant. \square

4. Hardness Results

An SMI (Stable Marriage with Incomplete Lists) instance is just like an HR instance, only that all firms have capacity 1. Hence, all of its stable matchings have the same size. An SMTI (Stable Marriage with Ties and Incomplete Lists) instance is an SMI instance except that the participants' preference lists are allowed to contain ties. For this problem, a pair (f, w) is a blocking pair of matching μ if (i) (f, w) is an acceptable pair; (ii) f has an opening or f strictly prefers w to its only employee under μ ; and (iii) w is unemployed or w strictly prefers f to her employer under μ . Once again, a matching is (weakly) stable if it has no blocking pairs. Unlike SMI instances, the stable matchings of an SMTI instance can have different sizes. In this section, we will show that certain kinds of SMTI instances can be encoded as HR+SN instances. This will allow us to translate hardness results known for max-SMTI, the problem of computing a maximum (cardinality) stable matching of an SMTI instance, to max-HR+SN.

Let I be an SMTI instance. Suppose the ties in the preference lists of I are broken arbitrarily to create the SMI instance I' . Clearly, every stable matching of I' is also a stable matching of I . The converse, however, is not true. Let us say that the ties in I are *consistent* if for every pair of participants q and q' , whenever q and q' appear in the preference lists of p and p' , q and q' are in a tie in the preference list of p if and only if they are also in a tie in the preference list of p' . In the next theorem, we show that when only the firms' preference lists contain ties and these ties are consistent, then the stable matchings of I can be retrieved from I' by considering the locally stable matchings of (I', G) , where G is constructed appropriately.

Theorem 5. *Let I be an SMTI instance where only the firms' preference lists contain ties, and the ties are consistent. Let I' be the SMI instance obtained by breaking the ties in the preference lists of I arbitrarily. Let G be a graph such that whenever two workers w and w' appear together in some firm's preference list in I , w and w' are adjacent if and only if they are not in a tie. Then the following are true:*

- (i) *Every stable matching of I is also a locally stable matching of (I', G) .*
- (ii) *Every locally stable matching μ' of (I', G) can be transformed into a stable matching μ of I such that $|\mu| \geq |\mu'|$ in time polynomial in the size of I . Consequently, every maximum locally stable matching of (I', G) can be transformed into a maximum stable matching of I of the same size in time polynomial in $|I|$.*

Proof. For (i), let μ be a stable matching of I . Suppose μ is not a locally stable matching of (I', G) , and (f, w) is one of its local blocking pairs with w' as a point of contact. Since f has capacity 1, w' is the only employee of f and f prefers w over w' . In order for (f, w) to not to be a blocking pair of μ in I , w and w' must be in a tie in f 's preference list in I . However, this cannot be the case— w and w' are adjacent in G and ties are consistent in I . Hence, μ cannot have a local blocking pair in (I', G) and must therefore be a locally stable matching of the instance.

For (ii), suppose μ' is a locally stable matching of (I', G) . Let (f, w) be a blocking pair of μ' in I . Without loss of generality, assume that w is the worker that f prefers the most among those that form a blocking pair with f . First, we note that f cannot be matched in μ' . Otherwise, if it is matched to some worker w_1 then f must strictly prefer w over w_1 so that w and w_1 are adjacent in G . In (I', G) , (f, w) has w_1 as a point of contact, implying that μ' cannot be a locally stable matching because (f, w) is a

local blocking pair. Since this is a contradiction, f has to be unmatched in μ' . Next, if w is unmatched in μ' , let $\mu'' = \mu' \cup \{(f, w)\}$; otherwise, let $\mu'' = \mu' - \{(\mu'(w), w)\} \cup \{(f, w)\}$. If μ'' has a local blocking pair, it will involve either f or w . By our choice of w , no worker will form a local blocking pair with f . If (f_1, w) is a local blocking pair of μ'' , then (f_1, w) must be a local blocking pair of μ' too—a contradiction. Hence, μ'' is still a locally stable matching of (I', G) .

What we have shown is that as long as a locally stable matching of (I', G) has a blocking pair with respect to I , the matching can be modified to a new matching (which no longer admits that blocking pair with respect to I) such that (i) its size stays the same or is larger by 1; and (ii) the modified matching is still a locally stable matching of (I', G) where one worker's employer improved while everyone else's stayed the same. If we keep applying this modification, at some point there will be no more worker whose employer can be improved. The locally stable matching of (I', G) obtained after this process will also be a stable matching of I .

Checking whether μ' is a stable matching of I , finding a blocking pair of μ' if it is not, and finding a worker that f prefers the most among those that form a blocking pair with f can be done in time polynomial in the size of I . Since the number of modifications cannot be more than $|F| \times |W|$, it follows that starting at a locally stable matching μ' of (I', G) , we can arrive at a stable matching μ of I such that $|\mu| \geq |\mu'|$ in time polynomial in the size of I .

Finally, we note from (i) that a maximum locally stable matching of (I', G) is at least as large as a maximum stable matching of I . Hence, if a maximum locally stable matching of (I', G) is *also* a stable matching of I , it must be a maximum stable matching of I . Our argument in the previous paragraphs shows that the last part of (ii) is true. \square

In the statement of Theorem 5, we simply described what edges should be in G : whenever two workers w and w' appear together in some firm's preference list in I , w and w' are adjacent if and only if they are not in a tie. That is, we are ambivalent about edges formed by workers that do not appear together in a firm's preference list in I since according to Proposition 4, the absence or presence of these edges in G have no effect on the set of locally stable matchings of (I, G) .

Consistent ties arise naturally when firms and/or workers derive their preference lists from a *master list* [13]. A master list of workers L_W is an ordering of all the workers that may or may not contain ties. Each firm's preference list contains all the workers acceptable to it and ranked in accordance with the master list. Thus, when w and w' is part of the preference list of a firm f , they are in a tie in f 's list if and only if they are in a tie in L_W . A master list of firms L_F is defined similarly, and each worker's preference list is obtained in the same way. Let SMTI-2ML denote the SMTI problem where both groups of participants derived their preference lists from a master list. The following result is known about SMTI-2ML.

Fact 1. (Irving et al. [13]) *Suppose that in the SMTI-2ML instance I , there are n firms and n workers. Determining if (I, G) has a stable matching of size n is NP-complete even if the ties occur in one master list only. The result holds even when (i) there is only one tie in that master list or (ii) all the ties are of length 2.*

We now translate this result to HR+SN.

Theorem 6. Suppose that in the HR+SN instance (I, G) , there are n firms and n workers, and each firm has capacity 1. Let K_z denote the complete graph on z vertices. Determining if I has a locally stable matching of size n is NP-complete, even when $G \cong K_n - K_{n'}$ where $n' < n$ or $G \cong K_n - F$ where F is a matching in K_n .

Proof. Let I' be an SMTI-2ML instance with n firms and n workers, and the ties occur in one master list only. Without loss of generality, assume that it is the firms' master list of workers that contains the ties. Create the HR+SN instance (I'', G) according to Theorem 5. If there is an efficient algorithm for determining if (I'', G) has a locally stable matching of size n , which is clearly a maximum locally stable matching of (I'', G) , then there is also an efficient algorithm for determining if I' has a stable matching of size n . However, Fact 1 states that the latter is an NP-complete problem. It follows that determining if (I'', G) has a locally stable matching of size n is also NP-complete. The form of G is based on Proposition 4, Fact 1 and Theorem 5. □

Next, we argue that max-HR+SN is NP-hard to approximate by appealing to the details of the following result of Halldórsson *et al.* [14] with regards to approximating max-SMTI.

Fact 2. (Halldórsson *et al.* [14]) It is NP-hard to approximate max-SMTI within a factor of $21/19 - \delta$ for any constant $\delta > 0$.

Theorem 7. It is NP-hard to approximate max-HR+SN within a factor of $21/19 - \delta$ for any constant $\delta > 0$.

Proof. To prove Fact 2, Halldórsson *et al.* [14] relied on a result of Dinur and Safra [15] about approximating a minimum vertex cover of a graph. Given a graph $H = (V, E)$, we now describe how they constructed the SMTI instance I_H . Note that we shall use firms in place of men and workers in place of women. For each vertex v_i of H , create three firms v_i^A, v_i^B, v_i^C , and three workers v_i^a, v_i^b, v_i^c . Thus, there are a total of $3|V|$ firms and $3|V|$ workers. Suppose v_i is adjacent to d vertices $v_{i_1}, v_{i_2}, \dots, v_{i_d}$. Here are the preference lists of the firms and workers corresponding to v_i :

$$\begin{array}{ll}
 v_i^A: & v_i^a & & & v_i^a: & v_i^B & v_{i_1}^C & \dots & v_{i_d}^C & v_i^A \\
 v_i^B: & (v_i^a & v_i^b) & & v_i^b: & v_i^B & v_i^C & & & \\
 v_i^C: & v_i^b & v_{i_1}^a & \dots & v_{i_d}^a & v_i^c & & & v_i^c: & v_i^C
 \end{array}$$

Clearly, I_H can be constructed from H in time polynomial in the size of H . Also, notice that the ties in I_H are consistent since v_i^a and v_i^b appear together in v_i^B 's preference list only. Using the reduction in Theorem 5, let (I'_H, G) be the HR+SN instance that corresponds to I_H . Let $s^+(I_H)$ and $s^+(I'_H, G)$ denote the sizes of a maximum stable matching in I_H and a maximum locally stable matching in (I'_H, G) respectively. According to Theorem 5, $s^+(I_H) = s^+(I'_H, G)$. Let $\beta^-(H)$ denote the size of a minimum vertex cover of H . They showed that $s^+(I_H) = 3|V| - \beta^-(H)$. Since $\beta^-(H)$ is NP-hard to approximate, it follows that $s^+(I_H)$ is also NP-hard to approximate. Since $s^+(I_H) = s^+(I'_H, G)$, and HR+SN instance (I'_H, G) can be constructed from I_H in polynomial time, we also have that $s^+(I'_H, G)$ is NP-hard to approximate. We refer readers to [14] for the derivation of the factor $21/19 - \delta$. □

5. Final Remarks

Theorem 7 provides a lower bound while Proposition 3 provides an upper bound to the factor of the best approximation algorithm for max-HR+SN. Can this gap be narrowed? We suspect that the answer is yes since the source of our hardness results, max-SMTI, has a number of $3/2$ -approximation algorithms [16–18]. An important strategy in these algorithms is to come up with a stable matching μ such that if μ^* is a maximum stable matching of the instance, there is no length-3 μ^* -alternating path in $\mu^* \oplus \mu$. In other words, in all μ^* -alternating paths in $\mu^* \oplus \mu$, the ratio of the number of μ^* -edges to the number of μ -edges is at most $3 : 2$. We end with the theorem below, which shows that for some HR+SN instances, choosing μ to be equal to a stable matching of the instance yields such a result.

Theorem 8. *In the HR+SN instance (I, G) , let W' be the set of workers that get matched in every stable matching of I . Suppose that in graph G , whenever $w_1 \in W'$ and $w_2 \in W - W'$ have a firm in common in their preference lists, w_1 and w_2 are adjacent in G . Let μ be a stable matching of I and $\hat{\mu}$ be a maximum locally stable matching of (I, G) . Then $|\hat{\mu}| \leq \frac{3}{2}|\mu|$.*

Proof. Recall the proof of Theorem 1. For now, assume all firms in I have capacity 1. Consider $\hat{\mu} \oplus \mu$. Suppose it has a length-3 $\hat{\mu}$ -alternating path: f_1, w_1, f_2, w_2 . This means that f_1 is unmatched in μ so w_1 must prefer f_2 over f_1 . Similarly, w_2 is unmatched in μ so f_2 must prefer w_1 over w_2 . Thus, (f_2, w_1) is a blocking pair of $\hat{\mu}$. Furthermore, $w_1 \in W'$ and $w_2 \in W - W'$ are adjacent in G so w_2 is a point of contact between f_2 and w_1 and (f_2, w_1) is a local blocking pair of $\hat{\mu}$ —a contradiction. It follows that $\hat{\mu} \oplus \mu$ has no length-3 $\hat{\mu}$ -alternating paths.

Now, suppose some firms in I have capacity greater than 1. Once more, construct $(I^{(1)}, G)$ from (I, G) . Let τ be a maximum locally stable matching of $(I^{(1)}, G)$. From the previous paragraph, $|\tau| \leq \frac{3}{2}|\mu^{(1)}|$. Thus, $|\hat{\mu}|/|\mu| = |\hat{\mu}^{(1)}|/|\mu^{(1)}| \leq |\tau|/|\mu^{(1)}| \leq 3/2$. \square

References

1. Arcaute, E.; Vassilvitskii, S. Social Networks and Stable Matchings in the Job Market. In Proceeding of WINE '09: the 5th Workshop on Internet and Network Economics, Rome, Italy, 14–18 December 2009; Springer: Berlin/Heidelberg, Germany, 2009; Volume 5929, *Lecture Notes in Computer Science*, pp. 220–231.
2. Hofer, M. Local Matching Dynamics in Social Networks. In Proceedings of ICALP '11: 39th International Colloquium on Automata, Languages and Programming, Zurich, Switzerland, 4–8 July 2011; Springer: Berlin/Heidelberg, Germany, 2011; Volume 6756, *Lecture Notes in Computer Science*, pp. 113–124.
3. Abraham, D.; Irving, R.; Kavitha, T.; Mehlhorn, K. Popular matchings. *SIAM J. Comput.* **2007**, *37*, 1030–1045.
4. Mestre, J. Weighted Popular Matchings. In Proceedings of ICALP '06: the 33rd International Colloquium on Automata, Languages and Programming, Venice, Italy, 9–16 July 2006; Springer: Berlin/Heidelberg, Germany, 2006; Volume 4051, *Lecture Notes in Computer Science*, pp. 715–726.

5. McCutchen, R. The Least-unpopularity-factor and Least-unpopularity-margin Criteria for Matching Problems with One-sided Preferences. In Proceedings of LATIN '08: the 8th Latin-American Theoretical Informatics Symposium, Rio De Janeiro, Brazil, 7–11 April 2008; Springer: Berlin/Heidelberg, Germany, 2008; Volume 4957, *Lecture Notes in Computer Science*, pp. 593–604.
6. Kavitha, T.; Nasre, M. Popular Matchings with Variable Job Capacities. In Proceedings of ISAAC '09: the 20th International Symposium on Algorithms and Computation, Honolulu, USA, 16–18 December 2009; Springer-Verlag: Berlin/Heidelberg, Germany, 2009; pp. 423–433.
7. Irving, R.; Kavitha, T.; Mehlhorn, K.; Michail, D.; Paluch, K. Rank-maximal matchings. *ACM Trans. Algorithms* **2006**, *2*, 602–610.
8. Biró, P.; Manlove, D.F.; Mittal, S. Size versus stability in the marriage problem. *Theor. Comput. Sci.* **2010**, *411*, 1828–1841.
9. Hoefler, M.; Wagner, L. Locally stable matching with general preferences. **2012**, arXiv:1207.1265v2.
10. Gale, D.; Shapley, L. College admissions and the stability of marriage. *Am. Math. Mon.* **1962**, *69*, 9–15.
11. Gale, D.; Sotomayor, M. Some remarks on the stable matching problem. *Discret. Appl. Math.* **1985**, *11*, 223–232.
12. Gusfield, D.; Irving, R. *The Stable Marriage Problem: Structure and Algorithms*; MIT Press: Cambridge, MA, USA, 1989.
13. Irving, R.; Manlove, D.; Scott, S. The stable marriage problem with master preference lists. *Discret. Appl. Math.* **2008**, *156*, 2959–2977.
14. Halldórsson, M.; Iwama, K.; Miyazaki, S.; Yanagisawa, H. Improved approximation of the stable marriage problem. *ACM Trans. Algorithms* **2007**, *3:3*, Article No. 30.
15. Dinur, I.; Safra, S. The Importance of Being Biased. In Proceedings of STOC '02: the 34th Symposium on the Theory of Computing, Montreal, Canada, 21–24 May 2002; ACM: New York, NY, USA, 2002; pp. 33–42.
16. McDermid, E. A $3/2$ -approximation Algorithm for General Stable Marriage. In Proceedings of ICALP '09: the 36th International Colloquium on Automata, Languages and Programming, Rhodes, Greece, 5–12 July 2009; Lecture Notes in Computer Science, pp. 689–700.
17. Király, Z. Better and simpler approximation algorithms for the stable marriage problem. *Algorithmica* **2011**, *60*, 3–20.
18. Paluch, K. Faster and simpler approximation of stable matchings. 2011. In Proceedings WAOA '11: the 9th Workshop on Approximation and Online Algorithms, Saarbrücken, Germany, 8–9 September 2011; Lecture Notes in Computer Science, pp. 176–187.