



Envy-Freeness and Relaxed Stability: Hardness and Approximation Algorithms

Prem Krishnaa¹, Girija Limaye¹, Meghana Nasre¹,
and Prajakta Nimbhorkar^{2,3}(✉)

¹ Indian Institute of Technology Madras, Chennai, India
premkrisnaa.jaganmohan@cohesity.com, {girija,meghana}@cse.iitm.ac.in

² Chennai Mathematical Institute, Chennai, India
prajakta@cmi.ac.in

³ UMI ReLaX, Chennai, India

Abstract. We consider the problem of matchings under two-sided preferences in the presence of maximum as well as minimum quota requirements for the agents. When there are no minimum quotas, *stability* is the de-facto notion of optimality. In the presence of minimum quotas, ensuring stability and simultaneously satisfying lower quotas is not an attainable goal in many instances.

To address this, a relaxation of stability known as *envy-freeness*, is proposed in literature. In our work, we thoroughly investigate envy-freeness from a computational view point. Our results show that computing envy-free matchings that match maximum number of agents is computationally hard and also hard to approximate up to a constant factor. Additionally, it is known that envy-free matchings satisfying lower-quotas may not exist. To circumvent these drawbacks, we propose a new notion called *relaxed stability*. We show that relaxed stable matchings are guaranteed to exist even in the presence of lower-quotas. Despite the computational intractability of finding a largest matching that is feasible and relaxed stable, we give an efficient algorithm that computes a constant factor approximation to this matching in terms of size.

Keywords: Matchings under preferences · Lower quota · Envy-freeness · Relaxed stability · Approximation

1 Introduction

Matching problems with two-sided preferences have been extensively investigated for matching markets where agents (hospitals/residents or colleges/students) have upper quotas that cannot be exceeded. Stability [6] is a widely accepted notion of optimality in this scenario. An allocation is said to be *stable* if no pair of agents has an incentive to deviate from it. However, the case when the agents

This work was partially supported by the grant CRG/2019/004757.

P. Krishnaa—Part of this work was done when the author was a student at IIT Madras.

© Springer Nature Switzerland AG 2020

T. Harks and M. Klimm (Eds.): SAGT 2020, LNCS 12283, pp. 193–208, 2020.

https://doi.org/10.1007/978-3-030-57980-7_13

have maximum as well as minimum quotas poses new challenges and there is still a want of satisfactory mechanisms that take minimum quotas into account. Practically, lower quotas are important, since it is natural for a hospital to require a minimum number of residents to run the hospital smoothly. Lower quotas are crucial in applications like course-allocation, and assigning teaching assistants (TAs) in academic institutions where a minimum guarantee is essential.

Ensuring stability while satisfying lower quotas may not be attainable always. On one hand, disregarding preferences in the interest of satisfying the lower quotas gives rise to social unfairness (for instance agents envying each other); on the other hand, too much emphasis on fairness can lead to *wastefulness* [5]. Hence, it is necessary to strike a balance between these three mutually conflicting goals – optimality with respect to preferences, feasibility for minimum quotas and minimizing wastefulness. The main contribution of this paper is to propose a mechanism to achieve this balance.

Envy-freeness [3, 5, 7, 11, 12] is a widely accepted notion for achieving fairness from a social perspective. Unfortunately, even envy-freeness and feasibility may not be simultaneously achievable. Whether feasible envy-free matchings exist can be answered efficiently by the characterization of Yokoi [20]. Fragiadakis et al. [5] explore strategyproofness and the trade-off between envy-freeness and wastefulness for a restricted setting of agent preferences. In our work, we thoroughly investigate envy-freeness from a computational view point. Our results show that computing a maximum size envy-free matching is computationally hard and such matchings can be wasteful. To circumvent these drawbacks, we propose a new notion called *relaxed stability*. We show that relaxed stable matchings are guaranteed to exist even in the presence of lower-quotas. Despite the computational intractability of finding a largest feasible relaxed stable matching, we give an efficient constant-factor approximation algorithm for it.

We state the problem formally in terms of a setting known as the HRLQ setting in literature. An HRLQ instance consists of a bipartite graph $G = (\mathcal{R} \cup \mathcal{H}, E)$, \mathcal{R} and \mathcal{H} being the sets of residents and hospitals respectively, and an edge $(r, h) \in E$ denotes that r and h are mutually acceptable. Each $h \in \mathcal{H}$ has an upper-quota $q^+(h)$ and a lower-quota $q^-(h)$, respectively denoting the maximum and minimum number of residents that can be assigned to h . Every vertex in $\mathcal{R} \cup \mathcal{H}$ ranks its neighbors in a strict order, referred to as its *preference list*. If a vertex a prefers its neighbor b_1 over b_2 , we denote it by $b_1 >_a b_2$.

A matching $M \subseteq E$ in G is an assignment of residents to hospitals such that each resident is matched to at most one hospital, and every hospital h is matched to at most $q^+(h)$ -many residents. Let $M(r)$ denote the hospital that r is matched to in M , and $M(h)$ denote the set of residents matched to h in M . We let $M(r) = \perp$ if r is unmatched in M , and \perp is considered as the least preferred choice of each $r \in \mathcal{R}$. We say that a hospital h is *under-subscribed* in M if $|M(h)| < q^+(h)$, is *fully-subscribed* if $|M(h)| = q^+(h)$ and is *deficient* if $|M(h)| < q^-(h)$. A matching is *feasible* for an HRLQ instance if no hospital is deficient in M . The goal in the HRLQ setting is to find a feasible matching M that is *optimal* with respect to the preference lists. The HRLQ problem is a generalization of the well-studied HR

problem (introduced by Gale and Shapley [6]) where there are no lower quotas. In the HR problem, *stability* is a de-facto notion of optimality and is defined by the absence of *blocking pairs*.

Definition 1 (Stable matchings). A pair $(r, h) \in E \setminus M$ is a *blocking pair* w.r.t. the matching M if $h >_r M(r)$ and h is either under-subscribed in M or there exists at least one resident $r' \in M(h)$ such that $r >_h r'$. A matching M is *stable* if there is no blocking pair w.r.t. M .

Existence of Stable Feasible Matchings:

Given an HRLQ instance, it is natural to ask “does the instance admit a stable feasible matching?” Unlike HR instances, an HRLQ instance may not admit a stable, feasible matching. Figure 1 shows an example. The stable matching $M_s = \{(r_1, h_1)\}$ is not feasible since h_2 is deficient in M_s , and the feasible matchings are not stable. The well-known Rural Hospitals Theorem [18] implies that the *number* of residents matched to a hospital is invariant across all stable matchings of the instance. Hence, for any HRLQ instance, either all stable matchings are feasible or all are infeasible. In light of the fact that stable and feasible matchings may not exist, relaxations of stability, like popularity and envy-freeness have been proposed in the literature [16, 17, 20]. Envy-freeness is defined by the absence of *envy-pairs*.

$$\begin{array}{ll} r_1 : h_1, h_2 & [0, 1] \quad h_1 : r_1, r_2 \\ r_2 : h_1 & [1, 1] \quad h_2 : r_1 \end{array}$$

Fig. 1. An HRLQ instance with no feasible and stable matching. Here $\mathcal{R} = \{r_1, r_2\}$, $\mathcal{H} = \{h_1, h_2\}$ and quotas are denoted as [lower-quota, upper-quota] pair preceding each hospital.

Definition 2 (Envy-free matchings). Given a matching M , a resident r has a *justified envy* (here onwards called *envy*) towards a matched resident r' , where $M(r') = h$ and $(r, h) \in E$ if $h >_r M(r)$ and $r >_h r'$. The pair (r, r') is an *envy-pair* w.r.t. M . A matching is *envy-free* if there is no envy-pair w.r.t. it.

Note that an envy-pair implies a blocking pair but the converse is not true and hence envy-freeness is a relaxation of stability. In the example in Fig. 1, the matching $\{(r_1, h_2)\}$ is envy-free and feasible, although not stable. Thus, envy-free matchings provide an alternative to stability in such instances. Envy-freeness is motivated by fairness from a social perspective. Importance of envy-free matchings has been recognized in the context of constrained matchings [3, 5, 7, 11, 12], and their structural properties have been investigated in [19].

Size of Envy-Free Matchings:

In terms of size, there is a sharp contrast between stable matchings in the HR setting and envy-free matchings in the HRLQ setting. While all the stable matchings in an HR instance have the same size, the envy-free matchings in an HRLQ instance may have significantly different sizes. For example,

$$\forall i \in [n], \quad r_i : h_1, h_2 \quad \begin{array}{l} [0, n] \quad h_1 : r_1, \dots, r_n \\ [1, 1] \quad h_2 : r_1, \dots, r_n \end{array}$$

Fig. 2. An HRLQ instance with two envy-free matchings of different sizes.

in Fig. 2, there are two envy-free matchings, $N_1 = \{(r_1, h_2)\}$ of size one and $N_n = \{(r_1, h_1), (r_2, h_1), \dots, (r_{n-1}, h_1), (r_n, h_2)\}$ of size n .

Shortcomings of Envy-Free Matchings: It is interesting to note that a feasible, envy-free matching itself may not exist – e.g. in Fig. 1, if both h_1, h_2 have a unit lower-quota, then the unique feasible matching is not envy-free. If a stable matching is not feasible in an HRLQ instance, *wastefulness* may be inevitable for attaining feasibility. A matching is *wasteful* if there exists a resident who prefers a hospital to her current assignment and that hospital has a vacant position [5]. Envy-free matchings can be significantly wasteful (e.g. the matching N_1 in Fig. 2). Therefore, it would be ideal to have a notion of optimality which is guaranteed to exist, is efficiently computable and avoids wastefulness.

Quest for a Better Optimality Criterion: We propose a new notion of *relaxed stability* which always exists for any HRLQ instance. We observe that in the presence of lower quotas, there can be at most $q^-(h)$ -many residents that are forced to be matched to h , even though they have higher preferred under-subscribed hospitals in their list. Our relaxation allows these forced residents to participate in blocking-pairs,¹ however, the matching is still stable when restricted to the remaining residents. We now make this formal below.

Definition 3 (Relaxed stable matchings). *A matching M is relaxed stable if, for every hospital h , at most $q^-(h)$ residents from $M(h)$ participate in blocking pairs and no unmatched resident participates in a blocking pair.*

In Fig. 1, the matching $\{(r_1, h_2), (r_2, h_1)\}$ (which was not envy-free) is feasible, relaxed stable and non-wasteful. We show that a feasible relaxed stable matching always exists in an HRLQ instance. However, computing a largest relaxed stable matching is NP-hard. We present an efficient algorithm that computes a match-

| | | |
|------------------|----------|------------------|
| $r_1 : h_1, h_3$ | $[0, 1]$ | $h_1 : r_1$ |
| $r_2 : h_2, h_3$ | $[0, 1]$ | $h_2 : r_2, r_3$ |
| $r_3 : h_2$ | $[1, 1]$ | $h_3 : r_1, r_2$ |

Fig. 3. An HRLQ instance with two relaxed stable matchings of different sizes, one larger than stable matching

ing that is at least as large as any stable matching in the instance, thus addressing wastefulness. In fact, a relaxed stable matching may be even larger than the stable matching in the instance. In the instance shown in Fig. 3, $M_s = \{(r_1, h_1), (r_2, h_2)\}$ is an infeasible stable matching. Matchings $M'_1 = \{(r_1, h_3), (r_2, h_2)\}$ and $M'_2 = \{(r_1, h_1), (r_2, h_3), (r_3, h_2)\}$ both are feasible, relaxed stable and M'_2 is larger than M_s . This is in contrast to maximum size envy-free matching which cannot be larger than a stable matching (see Sect. 2.2).

In the spirit of allowing blocking pairs, different notions have been proposed in [2, 9]. In [9], the goal is to compute a feasible matching with the least number

¹ Our initial idea was to allow them to participate in envy-pairs. We thank anonymous reviewer for suggesting this modification which is stricter than our earlier notion.

of blocking pairs or blocking residents; however, both these problems are NP-hard even for CL-restriction [9] (i.e., every hospital with positive lower-quota must rank every resident, and hence, every resident must rank every hospital with positive lower-quota) whereas a feasible relaxed stable matching can be efficiently computed. *Popular* matchings [17] allow blocking pairs to address feasibility and guaranteed existence in the HRLQ setting. However, there is no known bound on the number of blocking residents in a largest popular matching, whereas the number of blocking residents in a relaxed stable matching is at most the sum of lower quotas of all the hospitals. Lower quotas with constraints [4, 10] and in a model where hospitals can be closed [1] are investigated.

Our Contributions: We denote the problem of computing a maximum size feasible envy-free matching (respectively a maximum size feasible relaxed stable matching) as the MAXEFM (respectively the MAXRSM) problem. Throughout the paper, we assume that our input HRLQ instance admits a feasible matching. In the interest of space, proofs of Theorems and Lemmas marked with (\star) are deferred to the full-version [15].

Results on Envy-Freeness: We show that the MAXEFM problem is NP-hard, and is hard to approximate below a constant factor.

Theorem 1 (\star) . *The MAXEFM problem is NP-hard and cannot be approximated within a factor of $\frac{21}{19} - \epsilon$ for $\epsilon > 0$ unless $P = NP$ even when every hospital has a quota of at most one.*

In light of the above negative result, we turn our attention to the approximation and tractable special cases. In practice it is common to have incomplete preference lists and in many cases the preference lists of residents may also be of constant size. A matching M is a *maximal envy-free matching* if addition of an edge to M violates either the upper-quota or envy-freeness. Prior to our work, no size guarantee of a maximal envy-free matching was known. Let ℓ_1 and ℓ_2 be the length of the longest preference list of a resident and a hospital respectively.

Theorem 2. *A maximal envy-free matching is*

- (I) *an ℓ_1 -approximation of MAXEFM when hospital quotas are at most one.*
- (II) (\star) *an $(\ell_1 \cdot \ell_2)$ -approximation of MAXEFM when quotas are unrestricted.*

Next, we consider the HRLQ instances with the CL-restriction [9]. In contrast to the NP-hardness results in [9], the MAXEFM problem is tractable under the CL-restriction.

Theorem 3. *There is a simple linear-time algorithm for the MAXEFM problem for CL-restricted HRLQ instances.*

Results on Relaxed Stability: We prove that the MAXRSM problem is NP-hard and is also hard to approximate, but has a better approximation behavior than the MAXEFM problem.

Theorem 4. *The MAXRSM problem is NP-hard and cannot be approximated within a factor of $\frac{21}{19} - \epsilon$ for $\epsilon > 0$ unless $P = NP$ even when every hospital has a quota of at most one.*

We complement the above negative result with the following.

Theorem 5. *Any feasible HRLQ instance always admits a relaxed stable matching. Moreover, there is a polynomial-time algorithm that outputs a $\frac{3}{2}$ -approximation to the maximum size relaxed stable matching.*

We summarize our results in Table 1.

Table 1. Summary of our results

| Problem | Inapproximability | Approximation | Restricted settings |
|---------|---|--|---|
| MAXEFM | $(\frac{21}{19} - \epsilon)$ -inapproximability | $(\ell_1 \cdot \ell_2)$ -approximation | P-time for CL-restriction, ℓ_1 -approximation for 0/1 quotas |
| MAXRSM | $(\frac{21}{19} - \epsilon)$ -inapproximability | $\frac{3}{2}$ -approximation | – |

Organization of the Paper: Our algorithmic results for envy-free matchings and relaxed stable matchings are presented in Sect. 2 and in Sect. 3 respectively. The NP-hardness and inapproximability results are presented in Sect. 4.

2 Envy-Freeness: Algorithmic Results

In this section, we first focus on the approximation guarantee of maximal envy-free matchings and then present an efficient algorithm for the MAXEFM problem on the CL-restricted HRLQ instances.

2.1 Approximation to MAXEFM

A maximal envy-free matching can be efficiently computed; Krishnapriya et al. [16] present one such algorithm which extends a given envy-free matching. The results in [16] are empirical and no theoretical guarantee is known about the size of a maximal envy-free matching. Below we prove the guarantee for the instances where hospital quotas are at most one.

Proof (of Theorem 2(I)). Let M and OPT be respectively a maximal and a maximum size envy-free matching. Let R_{OPT} and R_M denote the set of residents matched in OPT and M respectively. Let X_1 be the set of residents matched in both M and OPT . Let X_2 be the set of residents matched in OPT but not

matched in M . Thus, $|R_{OPT}| = |X_1| + |X_2|$. Since $X_1 = R_{OPT} \cap R_M \subseteq R_M$, so $|X_1| \leq |R_M|$. Our goal is to show that $|X_2| \leq |R_M| \cdot (\ell_1 - 1)$. Recall that ℓ_1 is the length of the longest preference list of a resident. Once we establish that, it is immediate that a maximal envy-free matching is an ℓ_1 -approximation.

We show that for every resident $r \in X_2$ we can associate a unique hospital h_r such that h_r is unmatched in M and there exists a resident r' in the neighbourhood of h_r such that r' is matched in M . Denote the set of such hospitals as Y_2 . Note that due to the uniqueness assumption $|X_2| = |Y_2|$. Since each resident has a preference list of length at most ℓ_1 , any r' who is matched in M can have at most $\ell_1 - 1$ neighbouring hospitals which are unmatched in M . Thus $|X_2| = |Y_2| \leq |R_M| \cdot (\ell_1 - 1)$ which establishes the approximation guarantee. To finish the proof we show a unique hospital h_r with desired properties that can be associated with each $r \in X_2$. Let $r \in X_2$ such that $h = OPT(r)$. We have following two exhaustive cases.

Case 1: If h is unmatched in M , then due to maximality of M , there must exist a resident r' matched in M such that adding (r, h) causes envy to r' . Thus, h has a neighboring resident r' matched in M , and we let $h_r = h$.

Case 2: If h is matched in M , then since M and OPT are both envy-free, there must exist a path $\langle r, h, r_1, h_1, \dots, r_i, h_i \rangle$ such that $(r, h) \in OPT$, for each $k = 1, \dots, i$, we have $(r_k, h_k) \in OPT$, $(r_1, h) \in M$, for each $k = 2, \dots, i$, we have $(r_k, h_{k-1}) \in M$ and h_i is unmatched in M . Thus, h_i has a neighboring resident r_i matched in M , and we let $h_r = h_i$.

Uniqueness Guarantee: For any $r \in X_2$ such that case 1 applies, the associated h_i is unique since hospital quotas are at most 1. For two distinct $r, r' \in X_2$ such that case 2 applies for both, the paths mentioned above are disjoint since hospital quotas are at most 1, which guarantees uniqueness within case 2. The h_i associated in case 2 cannot be associated in case 1 to $OPT(h_i)$ since $OPT(h_i) = r_i \notin X_2$. This completes the proof of existence of the unique hospital. \square

2.2 Polynomial Time Algorithm for the CL-Restricted Instances

In this section, we consider the MAXEFM problem on CL-restricted HRLQ instances with general quotas. Recall that under CL-restriction [9], hospitals with positive lower-quota rank every resident and vice versa. It follows from the characterization of Yokoi [20] that every HRLQ instance with CL-restriction admits a feasible envy-free matching. We now present a simple modification to the standard Gale and Shapley algorithm [6] that computes a maximum size envy-free matching. We start with an empty matching M . Throughout the algorithm, we maintain two parameters:

- d : denotes the deficiency of the matching M , that is, the sum of deficiencies of all hospitals with positive lower-quota.
- k : the number of unmatched residents w.r.t. M .

In every iteration, an unmatched resident r who has not yet exhausted its preference list, proposes to the most preferred hospital h . If h is deficient w.r.t. M , h accepts r 's proposal. Otherwise, if h is under-subscribed w.r.t. M , h accepts the r 's proposal only if there are enough unmatched residents to satisfy the deficiency of the other hospitals, that is, $k > d$. If h is fully-subscribed, then h rejects the least preferred resident in $M(h) \cup \{r\}$. This process continues as long as some unmatched resident has not exhausted its preference list.

Algorithm 1: MAXEFM in CL-restricted HRLQ instances.

Input: An HRLQ instance $G = (\mathcal{R} \cup \mathcal{H}, E)$ with CL-restriction

Output: Maximum size envy-free matching

```

1 let  $M = \emptyset$ ;  $d = \sum_{h:q^-(h)>0} q^-(h)$ ;  $k = |\mathcal{R}|$ ;
2 while there is an unmatched resident  $r$  who has at least one hospital not yet
   proposed to do
3    $r$  proposes to the most preferred hospital  $h$ ;
4   if  $|M(h)| < q^-(h)$  then
5      $M = M \cup \{(r, h)\}$ ;
6     reduce  $d$  and  $k$  each by 1;
7   else
8     if  $|M(h)| == q^+(h)$  then
9       let  $r'$  be the least preferred resident in  $M(h) \cup \{r\}$ ;
10       $M(h) = M(h) \cup \{r\} \setminus \{r'\}$ ;
11     if  $|M(h)| < q^+(h)$  and  $k == d$  then
12       let  $r'$  be the least preferred resident in  $M(h) \cup \{r\}$ ;
13        $M(h) = M(h) \cup \{r\} \setminus \{r'\}$ ;
14     else
15       // we have  $|M(h)| < q^+(h)$  and  $k > d$ 
16        $M = M \cup \{(r, h)\}$ ;
17       reduce  $k$  by 1;
17 return  $M$ ;
```

Since the input instance is feasible, we start with $k \geq d$ and this inequality is maintained throughout the algorithm. If no resident is rejected due to $k = d$ in line 11, then our algorithm degenerates to the Gale and Shapley algorithm [6] and hence outputs a stable matching. It is straightforward to verify that Algorithm 1 runs in linear time in the size of the instance. Lemma 1 proves the correctness of our algorithm and establishes Theorem 3.

Lemma 1. *The matching M computed by Algorithm 1 is feasible and maximum size envy-free.*

Proof. We first prove that the output is feasible. Suppose not. Then at termination, $d > 0$, that is, there is at least one hospital h that is deficient w.r.t. M . It implies that $k \geq 1$. Thus there is some resident r unmatched w.r.t. M . Note

that r could not have been rejected by every hospital with positive lower-quota since h appears in the preference list of r , and h is deficient. This contradicts the termination of our algorithm and proves the feasibility of our matching.

Next, we prove that M is envy-free. Suppose for the sake of contradiction, M contains an envy-pair (r', r) such that $(r, h) \in M$ where $r' \succ_h r$ and $h \succ_{r'} M(r')$. This implies that r' must have proposed to h and h rejected r' . If h rejected r' because $|M(h)| = q^+(h)$, h is matched with better preferred residents than r' , a contradiction to the fact that $r' \succ_h r$. If h rejected r' because $k = d$, then there are two cases. Either r was matched to h when r' proposed to h . In this case, in line 11 our algorithm rejected the least preferred resident in $M(h)$. This contradicts that $r' \succ_h r$. Similarly if r proposed to h later, since $k = d$, the algorithm rejected the least preferred resident again contradicting the presence of any envy-pair.

Finally, we show that M is a maximum size envy-free matching. We have $k \geq d$ at the start of the algorithm. If during the algorithm, $k = d$ at some point, then at the end of the algorithm we have $k = d = 0$, implying that, we have an \mathcal{R} -perfect matching and hence the maximum size matching. Otherwise, $k > d$ at the end of the algorithm and then we output a stable matching which is maximum size envy-free by Lemma 2. \square

Lemma 2 (\star). *A stable matching when feasible, is an optimal solution of MAXEFM.*

Note that Algorithm 1 is similar to the ESDA algorithm presented in [5]. The ESDA algorithm needs a stricter assumption that the underlying graph is complete, whereas we assume the weaker CL-restriction. Moreover, only empirical results without theoretical guarantees on the size of the output matching are presented in [5].

3 Relaxed Stability: Algorithmic Results

In this section, we present Algorithm 2 that computes a relaxed stable matching in an HRLQ instance and prove that it gives a $\frac{3}{2}$ -approximation to MAXRSM. Furthermore, we show that the output of Algorithm 2 is at least as large as any stable matching in the instance (disregarding lower-quotas).

We say a feasible matching M_0 is *minimal* if, for any edge $e \in M_0$, $M_0 \setminus \{e\}$ is infeasible. Thus, if M_0 is minimal, then for every hospital h , $|M_0(h)| = q^-(h)$. Algorithm 2 begins by computing a feasible matching M_0 in the instance G disregarding the preferences of the residents and hospitals. Such a feasible matching can be computed by the standard reduction from bipartite matchings to flows with demands on edges [14]. Let $M = M_0$. We now associate levels with the residents – all residents matched in M are set to have level-0; all residents unmatched in M are assigned level-1. We now execute the Gale and Shapley resident proposing algorithm, with the modification that a hospital prefers any level-1 resident over any level-0 resident (irrespective of the preference list of h). Furthermore, if a level-0 resident becomes unmatched during the course of the proposals, then it

gets assigned a level-1 and it starts proposing from the beginning of its preference list. Amongst two residents of the same level, the hospital uses its preference list to order them. Our algorithm terminates when every resident is either matched or has exhausted its preference list when proposing hospitals at level-1. The two level idea is somewhat similar to the one used in Király [13] for stable matchings with ties and incomplete lists. It is clear that our algorithm runs in polynomial time since it only computes a feasible matching (using a reduction to flows) and executes a modification of Gale and Shapley algorithm. We prove the correctness of our algorithm below.

Algorithm 2: Algorithm to compute $\frac{3}{2}$ -approximation of MAXRSM

Input: Input: HRLQ instance $G = (\mathcal{R} \cup \mathcal{H}, E)$

Output: A relaxed stable matching that is a $\frac{3}{2}$ -approximation of MAXRSM

- 1 M_0 is a minimal feasible matching in G . Let $M = M_0$;
 - 2 For every matched resident r , set level of r to level-0;
 - 3 For every unmatched resident r , set level of r to level-1;
 - 4 **while** *there is an unmatched resident r who has not exhausted his preference list*
do
 - 5 r proposes to the most preferred hospital h to whom he has not yet proposed;
 - 6 **if** h *is under-subscribed* **then**
 - 7 $M = M \cup \{(r, h)\}$;
 - 8 **else**
 - 9 **if** $M(h)$ *has at least one level-0 resident r'* **then**
 - 10 $M = M \setminus \{(r', h)\} \cup \{(r, h)\}$;
 - 11 Set level of r' to level-1 and r' starts proposing from the beginning of his list;
 - 12 **else**
 - 13 h rejects the least preferred resident in $M(h) \cup \{r\}$;
 - 14 **Return** M ;
-

Remark 1. If r is unmatched in M then r is a level-1 resident and all the hospitals in r 's preference list are fully-subscribed with level-1 residents preferred over r .

Lemma 3. *Matching M output by Algorithm 2 is feasible and relaxed stable.*

Proof. We note that M_0 is feasible and since residents propose it is clear that for any hospital h , we have $|M(h)| \geq |M_0(h)| = q^-(h)$. Thus M is feasible.

To show relaxed stability, we claim that when the algorithm terminates, a resident at level-1 does not participate in a blocking pair. Whenever a level-1 resident r proposes to a hospital h , resident r always gets accepted except when h is fully-subscribed and all the residents matched to h are level-1 and are better preferred than r . When a matched level-1 resident r is rejected by a hospital h , h gets a better preferred resident than r . Thus, a level-1 resident

does not participate in a blocking pair. By Remark 1, an unmatched resident (being at level-1) does not participate in a blocking pair. Recall that all residents matched in M_0 are level-0 residents and M_0 is minimal. This implies that for every hospital h , at most $q^-(h)$ many residents assigned to h in M_0 participate in a blocking pair. We show that in M , the number of level-0 residents assigned to any hospital does not increase. To see this, if r is matched to h in M , but not matched to h in M_0 , it implies that either r was unmatched in M_0 or r was matched to some h' in M_0 . In either case r becomes level-1 when it gets assigned to h in M . Thus the number of level-0 residents assigned to any hospital h in M is at most $q^-(h)$, all of which can potentially participate in blocking pairs. This completes the proof that M is relaxed stable. \square

Lemma 4. *Matching M output by Algorithm 2 is a $\frac{3}{2}$ -approximation to the maximum size relaxed stable matching.*

Proof. Let OPT denote the maximum size relaxed stable matching in G . To prove the lemma we show that in $M \oplus OPT$ there does not exist any one-length as well as any three-length augmenting path. Suppose that (r, h) is a one-length augmenting path w.r.t. M in $M \oplus OPT$ implying that r is unmatched in M . Then by Remark 1, h is fully-subscribed - a contradiction that (r, h) is an augmenting path. Thus, there is no one-length augmenting path in $M \oplus OPT$.

For the three-length augmenting paths, we first convert the matchings M and OPT as one-to-one matchings, by making clones of the hospital. In particular we make $q^+(h)$ many copies of the hospital h for every h where the first $q^-(h)$ copies are called *lower-quota copies* and the $q^-(h) + 1$ to $q^+(h)$ copies are called *non lower-quota copies* of h . Let M_1 denote the one-to-one matching corresponding to M . To obtain M_1 , we assign every resident $r \in M(h)$ to a unique copy of h as follows: first, all the residents in $M(h)$ who participate in blocking pair w.r.t. M are assigned unique lower-quota copies of h arbitrarily. The remaining residents in $M(h)$ are assigned to the rest of the copies of h , ensuring all lower-quota copies get assigned some resident. We get OPT_1 from OPT in the same manner.

Now, suppose there exists a three-length augmenting path w.r.t. M which starts at an under-subscribed hospital, say h_j and ends at an unmatched resident in M . Since h_j is under-subscribed in M , and there is an augmenting path starting at h_j , it implies that there exists a copy h_j^d such that (i) h_j^d is matched in OPT_1 and unmatched in M_1 , say $OPT_1(h_j^d) = r_d$ and (ii) the resident r_d is matched in M_1 (otherwise there is a one-length augmenting path w.r.t. M_1 , which does not exist); let $M_1(r_d) = h_i^c$, and (iii) the copy h_i^c is matched in OPT_1 and $OPT_1(h_i^c) = r_c$ is unmatched in M_1 (else the claimed three-length augmenting path does not exist).

We first note that h_i^c and h_j^d are not copies of the same hospital, that is, $i \neq j$, otherwise there is a one-length augmenting path (r_c, h_i) w.r.t. M . Since r_c is unmatched, by Remark 1, r_d is a level-1 resident and $r_d >_{h_i} r_c$. Thus, r_d proposed to hospitals from the beginning of its preference list. Since h_j is under-subscribed, it must be the case that $h_i >_{r_d} h_j$. Thus, (r_d, h_i) is a blocking pair w.r.t. OPT . By the construction of OPT_1 from OPT , we must have assigned r_d

to a lower-quota copy of h_j . However, copy h_j^d is a non lower-quota copy, since it is unassigned in M_1 , a contradiction. Thus, the claimed three-length augmenting path does not exist. \square

We note that our analysis is tight [15]. We now show that the matching M computed by Algorithm 2 is at least as large as any stable matching.

Lemma 5 (\star). *A resident matched in a stable matching M_s is also matched in M output by Algorithm 2. Hence $|M| \geq |M_s|$.*

Proof (Sketch). Suppose there exists a resident r matched in M_s to hospital h but unmatched in M . We start constructing a path starting at r using edges from M_s and M alternately. We show that such a path can neither terminate at a resident nor at a hospital and hence cannot exist. Thus, every resident matched in M_s is matched in M and hence $|M| \geq |M_s|$. \square

4 Hardness Results

In this section we give an overview of the techniques used in proving the hardness and inapproximability results. Theorem 1 and Theorem 4 are proved using suitable reductions from the Minimum Vertex Cover (MVC) problem. We present the proof for Theorem 4 below.

Proof (of Theorem 4). Given a graph $G = (V, E)$, which is an instance of the MVC problem, we construct an instance G' of the MAXRSM problem. Corresponding to each vertex v_i in G , G' contains a gadget with three residents r_1^i, r_2^i, r_3^i , and three hospitals h_1^i, h_2^i, h_3^i . All hospitals have an upper-quota of 1 and h_3^i has a lower-quota of 1. Assume that the vertex v_i has d neighbors in G , namely v_{j_1}, \dots, v_{j_d} . The preference lists of the residents and hospitals are shown in Fig. 4. We impose an arbitrary but fixed ordering on the vertices which is used as a strict ordering of neighbors in the preference lists of resident r_1^i and hospital h_2^i in G' . Note that G' has $N = 3|V|$ residents and hospitals.

| | |
|--|---|
| $r_1^i : h_3^i, h_2^{j_1}, h_2^{j_2}, \dots, h_2^{j_d}, h_1^i$ | $[0, 1] h_1^i : r_1^i$ |
| $r_2^i : h_2^i, h_3^i$ | $[0, 1] h_2^i : r_2^i, r_1^{j_1}, r_2^{j_2}, \dots, r_2^{j_d}, r_3^i$ |
| $r_3^i : h_2^i$ | $[1, 1] h_3^i : r_2^i, r_1^i$ |

Fig. 4. Preferences of agents corresponding to a vertex v_i in G .

Lemma 6. *If $VC(G)$ denotes a minimum vertex cover of G and $OPT(G')$ denotes a maximum size relaxed stable matching in G' , then $|OPT(G')| = 3|V| - |VC(G)|$.*

Proof. We first prove that $|OPT(G')| \geq 3|V| - |VC(G)|$. Given a minimum vertex cover $VC(G)$ of G we construct a relaxed stable matching M for G' as follows. $M = \{(r_1^i, h_3^i), (r_2^i, h_2^i) \mid v_i \in VC(G)\} \cup \{(r_1^i, h_1^i), (r_2^i, h_3^i), (r_3^i, h_2^i) \mid v_i \notin VC(G)\}$. Thus, $|OPT(G')| \geq |M| = 2|VC(G)| + 3(|V| - |VC(G)|) = 3|V| - |VC(G)|$.

Claim. M is relaxed stable in G' .

Proof. When $v_i \in VC(G)$, residents r_1^i and r_2^i both are matched to their top choice hospitals and hospital h_2^i is matched to its top choice resident r_2^i . Thus, when $v_i \in VC(G)$, no resident from the i -th gadget participates in a blocking pair. When $v_i \notin VC(G)$, hospitals h_1^i and h_3^i are matched to their top choice residents and we ignore blocking pair (r_2^i, h_2^i) because r_2^i is matched to a lower-quota hospital h_3^i , thus there is no blocking pair within the gadget for $v_i \notin VC(G)$. Now suppose that there is a blocking pair (r_1^i, h_2^j) for some j such that $(v_i, v_j) \in E$. Note that either v_i or v_j is in $VC(G)$. If $v_i \in VC(G)$, r_1^i is matched to its top choice hospital h_3^i , thus cannot participate in a blocking pair. If $v_i \notin VC(G)$, it implies that $v_j \in VC(G)$. Then for v_j 's gadget, h_2^j is matched to its top choice r_2^j , thus cannot form a blocking pair. \square

Now we prove that $OPT(G') \leq 3|V| - |VC(G)|$. Let $M = OPT(G')$ be a maximum size relaxed stable matching in G' . Consider a vertex $v_i \in V$ and the corresponding residents and hospitals in G' . Refer Fig. 5 for the possible patterns caused by v_i . Hospital h_3^i must be matched to either resident r_1^i (Pattern 1) or resident r_2^i (Pattern 2 to Pattern 7). If $(r_1^i, h_3^i) \in M$, then the resident r_2^i must be matched to a higher preferred hospital h_2^i in M . If $(r_2^i, h_3^i) \in M$ then h_2^i may be matched with either r_3^i or r_1^j of some neighbour v_j or may be left unmatched. Similarly, r_1^i can either be matched to h_1^i or h_2^j of some neighbour v_j . This leads to 6 combinations as shown in Fig. 5b to Fig. 5g.

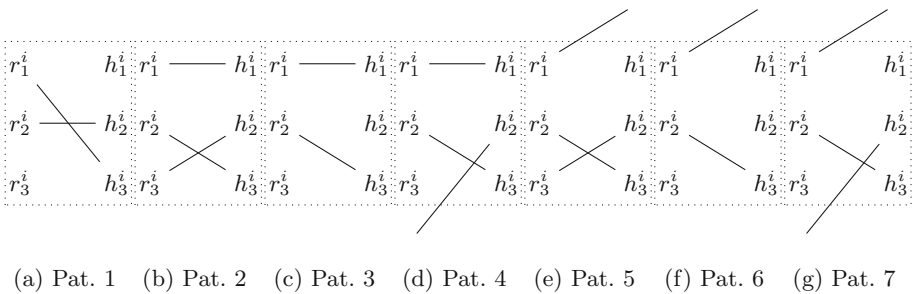


Fig. 5. Seven patterns possibly caused by vertex v_i

Claim. A vertex cannot cause pattern 5.

Proof. Assume for the sake of contradiction that a vertex v_i causes pattern 5. Then, there must exist a vertex v_j adjacent to v_i such that v_j causes either pattern 4 or pattern 7.

Case 1: If vertex v_j causes pattern 4, then (r_1^j, h_2^i) form a blocking pair, a contradiction.

Case 2: If vertex v_j causes pattern 7, then there must exist vertices v_{j+1}, \dots, v_t such that there are following edges in G : $(v_i, v_j), (v_j, v_{j+1}), (v_{j+1}, v_{j+2}), \dots, (v_{t-1}, v_t)$ and vertices v_j to v_{t-1} cause pattern 7 and v_t causes pattern 4. See Fig. 6. In the vertex ordering, we must have $v_{j+1} > v_i$ otherwise (r_1^j, h_2^i) form a blocking pair. But, since h_2^j is matched to r_1^i , $v_{j+2} > v_j$. Continuing this way, $v_t > v_{t-2}$ but this causes (r_1^t, h_2^{t-1}) form a blocking pair. Thus, the claimed set of edges cannot exist. \square




Fig. 6. Pattern combination that is not relaxed stable if v_i causes pattern 5

Claim. A vertex cannot cause pattern 3 or pattern 6 or pattern 4.

Proof. In pattern 3 and 6, r_3^i participates in a blocking pair (r_3^i, h_2^i) , contradicting that M is relaxed stable. If a vertex v_i causes pattern 4, then there exists a set of t vertices v_{i+1}, \dots, v_{i+t} such that for $0 \leq k < t$, (v_{i+k}, v_{i+k+1}) is an edge in G and v_{i+t} causes pattern 6. But, since pattern 6 cannot occur, pattern 4 cannot occur. \square

Thus, a vertex can cause either pattern 1 or 2 and thus match all the residents and hospitals within its own gadget or pattern 7 and match r_1 and h_2 outside its own gadget. Accordingly there are following cases.

Case 1: A vertex causing pattern 7 contributes size 1 for (r_2^i, h_3^i) edge and 0.5 each for two edges matched to another vertex causing pattern 7, contributing an average matching size of 2.

Case 2: It is clear that a vertex causing pattern 1 or 2 contributes to matching size of 2 or 3 respectively.

Vertex Cover C of G Corresponding to M : Using M , we now construct the set C of vertices in G which constitute a vertex cover of G . If v_i causes pattern 2, we do not include it in the C ; Otherwise, we include it. We prove that C is a vertex cover. Suppose not, then there exists an edge (v_i, v_j) such that both v_i and v_j cause pattern 2. But, this means that (r_1^i, h_2^j) and (r_1^j, h_2^i) form a blocking pair, a contradiction since M is relaxed stable. Now, it is easy to see that $|OPT(G')| = 2|C| + 3(|V| - |C|) = 3|V| - |C|$. Thus, $VC(G) \leq |C| = 3|V| - |OPT(G')|$. This completes the proof of the lemma. \square

The rest of our proof is similar to the approach of Halldórsson et al. [8] to prove inapproximability of the stable matchings with ties and incomplete lists; however our gadgets above are entirely different. Lemma 7 is analogous to Theorem 3.2 and Corollary 3.4 from [8] and its proof can be reproduced in a similar way [15].

Lemma 7. *It is NP-hard to approximate the MAXRSM problem within a factor of $\frac{21}{19} - \epsilon$, for any constant $\epsilon > 0$, even when the quotas of all hospitals are either 0 or 1.*

Acknowledgement. We thank the anonymous reviewers for their useful comments which has improved the presentation of the paper.

References

1. Biró, P., Fleiner, T., Irving, R.W., Manlove, D.: The college admissions problem with lower and common quotas. *Theoret. Comput. Sci.* **411**(34–36), 3136–3153 (2010). <https://doi.org/10.1016/j.tcs.2010.05.005>
2. Biró, P., Manlove, D., Mittal, S.: Size versus stability in the marriage problem. *Theoret. Comput. Sci.* **411**(16–18), 1828–1841 (2010). <https://doi.org/10.1016/j.tcs.2010.02.003>
3. Ehlers, L., Hafalir, I.E., Yenmez, M.B., Yildirim, M.A.: School choice with controlled choice constraints: hard bounds versus soft bounds. *J. Econ. Theory* **153**, 648–683 (2014). <https://doi.org/10.1016/j.jet.2014.03.004>
4. Fleiner, T., Kamiyama, N.: A matroid approach to stable matchings with lower quotas. *Math. Oper. Res.* **41**(2), 734–744 (2016). <https://doi.org/10.1287/moor.2015.0751>
5. Fragiadakis, D., Iwasaki, A., Troyan, P., Ueda, S., Yokoo, M.: Strategyproof matching with minimum quotas. *ACM Trans. Econ. Comput.* **4**(1), 6:1–6:40 (2015). <http://doi.acm.org/10.1145/2841226>
6. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. *Am. Math. Mon.* **69**(1), 9–15 (1962). <http://www.jstor.org/stable/2312726>
7. Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., Yokoo, M.: Strategyproof matching with regional minimum and maximum quotas. *Artif. Intell.* **235**, 40–57 (2016). <https://doi.org/10.1016/j.artint.2016.02.002>
8. Halldórsson, M.M., Iwama, K., Miyazaki, S., Yanagisawa, H.: Improved approximation results for the stable marriage problem. *ACM Trans. Algorithms* **3**(3), 30 (2007). <https://doi.org/10.1145/1273340.1273346>
9. Hamada, K., Iwama, K., Miyazaki, S.: The hospitals/residents problem with lower quotas. *Algorithmica* **74**(1), 440–465 (2016). <https://doi.org/10.1007/s00453-014-9951-z>
10. Huang, C.: Classified stable matching. In: *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010*, pp. 1235–1253 (2010). <https://doi.org/10.1137/1.9781611973075.99>
11. Kamada, Y., Kojima, F.: Efficient matching under distributional constraints: theory and applications. *Am. Econ. Rev.* **105**(1), 67–99 (2015). <https://www.aeaweb.org/articles?id=10.1257/aer.20101552>
12. Kamada, Y., Kojima, F.: Stability concepts in matching under distributional constraints. *J. Econ. Theory* **168**, 107–142 (2017). <https://doi.org/10.1016/j.jet.2016.12.006>

13. Király, Z.: Linear time local approximation algorithm for maximum stable marriage. *Algorithms* **6**(3), 471–484 (2013). <https://doi.org/10.3390/a6030471>
14. Kleinberg, J., Tardos, E.: *Algorithm Design*. Addison-Wesley Longman Publishing Co. Inc., Boston (2005)
15. Krishnaa, P., Limaye, G., Nasre, M., Nimbhorkar, P.: Envy-freeness and relaxed stability under lower quotas. CoRR abs/1910.07159 (2019). <http://arxiv.org/abs/1910.07159>
16. Krishnapriya, A.M., Nasre, M., Nimbhorkar, P., Rawat, A.: How good are popular matchings? In: 17th International Symposium on Experimental Algorithms, SEA 2018, pp. 9:1–9:14 (2018). <https://doi.org/10.4230/LIPIcs.SEA.2018.9>
17. Nasre, M., Nimbhorkar, P.: Popular matchings with lower quotas. In: 37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, pp. 44:1–44:15 (2017). <https://doi.org/10.4230/LIPIcs.FSTTCS.2017.44>
18. Roth, A.E.: On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica* **54**(2), 425–427 (1986). <http://www.jstor.org/stable/1913160>
19. Wu, Q., Roth, A.E.: The lattice of envy-free matchings. *Games Econ. Behav.* **109**, 201–211 (2018). <https://doi.org/10.1016/j.geb.2017.12.016>
20. Yokoi, Y.: Envy-free matchings with lower quotas. *Algorithmica* **82**(2), 188–211 (2020). <https://doi.org/10.1007/s00453-018-0493-7>