

The Hospitals/Residents Problem with Quota Lower Bounds

Koki Hamada¹, Kazuo Iwama^{2*}, and Shuichi Miyazaki^{3**}

¹ NTT Information Sharing Platform Laboratories, NTT Corporation
hamada.koki@lab.ntt.co.jp

² Graduate School of Informatics, Kyoto University
iwama@kuis.kyoto-u.ac.jp

³ Academic Center for Computing and Media Studies, Kyoto University
shuichi@media.kyoto-u.ac.jp

Abstract. The Hospitals/Residents problem is a many-to-one extension of the stable marriage problem. In its instance, each hospital specifies a quota, i.e., an upper bound on the number of positions it provides. It is well-known that in any instance, there exists at least one stable matching, and finding one can be done in polynomial time. In this paper, we consider an extension in which each hospital specifies not only an upper bound but also a *lower* bound on its number of positions. In this setting, there can be instances that admit no stable matching, but the problem of asking if there is a stable matching is solvable in polynomial time. In case there is no stable matching, we consider the problem of finding a matching that is “as stable as possible”, namely, a matching with a minimum number of blocking pairs. We show that this problem is hard to approximate within the ratio of $(|H| + |R|)^{1-\epsilon}$ for any positive constant ϵ where H and R are the sets of hospitals and residents, respectively. We tackle this hardness from two different angles. First, we give an exponential-time exact algorithm for a special case where all the upper bound quotas are one. This algorithm runs in time $O(t^2(|H|(|R| + t))^{t+1})$ for instances whose optimal cost is t . Second, we consider another measure for optimization criteria, i.e., the number of residents who are involved in blocking pairs. We show that this problem is still NP-hard but has a polynomial-time $\sqrt{|R|}$ -approximation algorithm.

1 Introduction

In the *stable marriage problem* [10], we are given sets of men and women, and each person’s preference list that orders the members of the other sex according to his/her preference. The question is to find a *stable matching*, that is, a matching containing no pair of man and woman who prefer each other to their partners. Such a pair is called a *blocking pair*. Gale and Shapley proved that any

* Supported by KAKENHI 22240001.

** Supported by KAKENHI 20700009.

instance admits at least one stable matching, and gave an algorithm to find one, known as the Gale-Shapley algorithm.

They also proposed a many-to-one extension of the stable marriage problem, which is currently known as the *Hospitals/Residents problem* (*HR* for short) [10]. In *HR*, the two sets corresponding to men and women are residents and hospitals. Each hospital specifies its *quota*, which means that it can accept at most this number of residents. Hence in a feasible matching, the number of residents assigned to each hospital is up to its quota. Most properties of the stable marriage problem carry over to *HR*, e.g., any instance admits a stable matching, and we can find one by the appropriately modified Gale-Shapley algorithm. As the name of *HR* suggests, it has real-world applications in assigning residents to hospitals in many countries, known as *NRMP* in the U.S. [12], *CaRMS* in Canada [7], and *SFAS* in Scotland [17]. Along with these applications and due to special requirements in reality, several useful extensions have been proposed, such as *HR* with couples [24, 23, 3, 22], and the Student-Project Allocation problem [2].

In this paper, we study another extension of *HR* where each hospital declares not only an upper bound but also a *lower* bound on the number of residents it accepts. Consequently, a feasible matching must satisfy the condition that the number of residents assigned to each hospital is between its upper and lower bound quotas. This restriction seems quite relevant in several situations. For example, shortage of doctors in hospitals in rural area is a critical issue; it is sometimes natural to guarantee some number of residents for such hospitals in the residents-hospitals matching. Also, when determining supervisors of students in universities, it is quite common to consider that the number of students assigned to each professor should be somehow balanced, which can be achieved again by specifying both upper and lower bounds on the number of students accepted by each professor. We call this problem *HR with Minimum Quota* (*HRMQ* for short).

The notion of minimum quota was first raised in [13] and followed by [5, 16] (see “Related Work” below). In this paper, we are interested in a most natural question, i.e., how to obtain “good” matchings in this new setting. In *HRMQ*, stable matchings do not always exist. However, it is easy to decide whether or not there is a stable matching for a given instance, since in *HR* the number of students a specific hospital h receives is identical for any stable matching (this is a part of the well-known *Rural Hospitals Theorem* [11]). Namely, if this number satisfies the upper and lower bound conditions of all the hospitals, it is a feasible (and stable) matching, and otherwise, no stable matching exists. In case there is no stable matching, it is natural to seek for a matching “as stable as possible”.

Our Contributions. We first consider the problem of minimizing the number of blocking pairs, which is quite popular in the literature (e.g., [20, 1, 6, 15]). As shown in Sec. 2, it seems that the introduction of the quota lower bound intrinsically increases the difficulty of the problem. Actually, we show that this problem is NP-hard and cannot be approximated within a factor of $(|H| + |R|)^{1-\varepsilon}$ for any positive constant ε unless $P=NP$, where H and R denote the sets of hospitals and residents, respectively. This inapproximability result holds even if

all the preference lists are complete, all the hospitals have the same preference list, (e.g., determined by scores of exams and known as the *master list* [18]), and all the hospitals have upper bound quota of one. On the positive side, we give a polynomial-time $(|H| + |R|)$ -approximation algorithm, which shows that the above inapproximability result is almost tight.

We then tackle this hardness from two different angles. First, we restrict ourselves to instances where upper bound quotas of all the hospitals are one, which correspond to the marriage case and are still hard to approximate as shown above. We give an exponential-time exact algorithm which runs in time $O(t^2(|H|(|R| + t))^{t+1})$ for instances whose optimal cost is t . Note that this is a polynomial-time algorithm when the optimal cost is constant. Second, we go back to the original many-to-one case, and consider another measure for optimization criteria, i.e., the number of residents who are involved in blocking pairs. We show that this problem is still NP-hard, but give a quadratic improvement, i.e., we give a polynomial-time $\sqrt{|R|}$ -approximation algorithm. We also give an instance showing that our analysis is tight up to a constant factor. Furthermore, we show that if our problem has a constant approximation factor, then the Dense k -Subgraph Problem (DkS) has a constant approximation factor also. Note that the best known approximation factor of DkS has long been $|V|^{1/3}$ [21] in spite of extensive studies, and was recently improved to $|V|^{1/4+\epsilon}$ for an arbitrary positive constant ϵ [4]. The reduction is somewhat tricky, which is done through the third problem, called the Minimum Coverage Problem (MinC), and exploits the best approximation algorithm for DkS. MinC is relatively less studied and only the NP-hardness is known for its complexity [25]. As a byproduct, our proof gives a similar hardness for MinC, which is of independent interest.

Because of the space restriction, most of the proofs are omitted. They are included in the full version [14].

Related Work. Biró, et al. [5] also considers HR with quota lower bounds. In contrast to our model, which requires to satisfy the lower bound quota of all the hospitals, their model allows some hospitals to be closed, i.e., to receive no residents. Huang [16] considers *classified stable matchings*, in which not only individual hospitals but also selected sets of hospitals declare quota upper and lower bounds. He proved a dichotomy theorem for the problem of deciding the existence of a stable matching, in terms of the structural property of the family of the sets of hospitals that declare quota bounds.

2 Preliminaries

An instance of the *Hospitals/Residents Problem with Minimum Quota* (HRMQ for short) consists of the set R of residents and the set H of hospitals. Each hospital h has lower and upper bounds of quota, p and q ($p \leq q$), respectively. We sometimes say that the quota of h is $[p, q]$, or h is a $[p, q]$ -hospital. For simplicity, we also write the name of a hospital with its quota bounds, such as $h[p, q]$. Each member (resident or hospital) has a preference list that orders a subset of the members of the other party.

A *matching* is an assignment of residents to hospitals (possibly, leaving some residents unassigned), where matched residents and hospitals are in the preference list of each other. Let $M(r)$ be the hospital to which resident r is assigned under a matching M (if it exists), and $M(h)$ be the set of residents assigned to hospital h . A *feasible matching* is a matching such that $p \leq |M(h)| \leq q$ for each hospital $h[p, q]$. We may sometimes call a feasible matching simply a matching when there is no fear of confusion. For a matching M and a hospital $h[p, q]$, we say that h is *full* if $|M(h)| = q$ and that h is *under-subscribed* if $|M(h)| < q$.

For a matching M , we say that a pair comprising a resident r and a hospital h who include each other in the list forms a *blocking pair* for M if the following two conditions are met: (i) r is either unassigned or prefers h to $M(r)$, and (ii) h is under-subscribed or prefers r to one of the residents in $M(h)$. We say that r is a *blocking resident* for M if r is involved in a blocking pair for M .

Minimum-Blocking-Pair HRMQ (*Min-BP HRMQ* for short) is the problem of finding a feasible matching with the minimum number of blocking pairs. *Min-BP 1ML-HRMQ* (“1ML” standing for “1 Master List”) is the restriction of Min-BP HRMQ so that preference lists of all the hospitals are identical. *0-1 Min-BP HRMQ* is the restriction of Min-BP HRMQ where a quota bound of each hospital is either $[0, 1]$ or $[1, 1]$. *0-1 Min-BP 1ML-HRMQ* is Min-BP HRMQ with both “1ML” and “0-1” restrictions. *Minimum-Blocking-Resident HRMQ* (*Min-BR HRMQ* for short) is the problem of finding a feasible matching with the minimum number of blocking residents. *Min-BR 1ML-HRMQ*, *0-1 Min-BR HRMQ*, and *0-1 Min-BR 1ML-HRMQ* are defined similarly.

We assume without loss of generality that the number of residents is at least the sum of the lower bound quotas of all the hospitals. Also, in this paper we impose the following restriction $\mathcal{Z}$ to guarantee existence of a feasible solution: every resident’s list includes all hospitals with positive quota lower bounds, and such hospitals’ lists include all the residents. (We remark in Sec. 5 that allowing arbitrarily incomplete preference lists makes the problem extremely hard.)

We say that an algorithm A is an $r(n)$ -approximation algorithm if it satisfies $A(x)/opt(x) \leq r(n)$ for any instance x of size n , where $opt(x)$ and $A(x)$ are the costs (i.e., the number of blocking pairs in the case of Min-BP HRMQ) of the optimal and the algorithm’s solutions, respectively.

As a starting example, consider n residents and $n + 1$ hospitals, whose preference lists and quota bounds are as follows. Here, “ $\dots$ ” in the residents’ preference lists denotes an arbitrary order of the remaining hospitals.

$$\begin{array}{llll}
r_1 : h_1 & h_{n+1} & \dots & \\
r_2 : h_1 & h_2 & h_n & \dots \\
r_3 : h_2 & h_1 & h_3 & \dots \\
r_4 : h_3 & h_1 & h_4 & \dots \\
\vdots & & & \\
r_i : h_{i-1} & h_1 & h_i & \dots \\
\vdots & & & \\
r_n : h_{n-1} & h_1 & h_n & \dots
\end{array}
\qquad
\begin{array}{l}
h_1[0, 1] : r_1 r_2 \dots r_n \\
h_2[1, 1] : r_1 r_2 \dots r_n \\
\vdots \\
h_n[1, 1] : r_1 r_2 \dots r_n \\
h_{n+1}[1, 1] : r_1 r_2 \dots r_n
\end{array}$$

Note that we have n $[1, 1]$ -hospitals all of which have to be filled by the n residents. Therefore, let us modify the instance by removing the $[0, 1]$ -hospital h_1 and apply the Gale-Shapley algorithm (see e.g., [12] for the Gale-Shapley algorithm; in this paper it is always the residents-oriented version, namely, residents make and hospitals receive proposals). Then the resulting matching is $M_1 = \{(r_1, h_{n+1}), (r_2, h_2), (r_3, h_3), \dots, (r_n, h_n)\}$, which contains at least n blocking pairs (between h_1 and every resident). However, the matching $M_2 = \{(r_1, h_{n+1}), (r_2, h_n), (r_3, h_2), (r_4, h_3), \dots, (r_n, h_{n-1})\}$ contains only three blocking pairs (r_1, h_1) , (r_2, h_1) , and (r_2, h_2) .

3 Minimum-Blocking-Pair HRMQ

In this section, we give both approximability and inapproximability results. For the latter, we prove a strong inapproximability result for the restricted subclass, as mentioned in Sec. 1. On the other hand, we can show that this inapproximability result is almost tight by providing an approximation algorithm for the general class.

Theorem 1. *For any positive constant ε , there is no polynomial-time $(|H| + |R|)^{1-\varepsilon}$ -approximation algorithm for 0-1 Min-BP 1ML-HRMQ unless $P=NP$, even if all the preference lists are complete. (Proof is omitted. See [14].)*

Theorem 2. *There is a polynomial-time $(|H| + |R|)$ -approximation algorithm for Min-BP HRMQ.*

Proof. Here we give only a sketch. See [14] for the complete proof. The following simple algorithm (Algorithm I) achieves an approximation ratio of $|H| + |R|$: Given an instance I of Min-BP HRMQ, consider it as an instance of HR by ignoring quota lower bounds. Then, apply the Gale-Shapley algorithm to I and obtain a matching M . In M , define a *deficiency of a hospital h_i* $[p_i, q_i]$ to be $\max\{p_i - x_i, 0\}$ where x_i is the number of residents assigned to h_i by M . We then move residents arbitrarily from hospitals with surplus to the hospitals with positive deficiencies (but so as not to create new positive deficiency) to fill all the deficiencies. This is possible because of the restriction $\mathcal{Z}$.

Let k be the sum of the deficiencies over all the hospitals. Then, k residents are moved by the above procedure. We can show that at most $|H| + |R|$ new blocking pairs can arise per resident movement and hence at most $k(|H| + |R|)$ in total. On the other hand, we can prove that if there are k deficiencies in M , an optimal solution contains at least k blocking pairs. $\square$

3.1 Exponential-Time Exact Algorithm

In this section we consider only the cases where quota bounds are $[0, 1]$ or $[1, 1]$, as in the example given in the previous section. Recall that the problem is still hard to approximate, and our goal here is to design nontrivial exponential-time algorithms by using the parameter t denoting the cost of an optimal solution. Probably, a natural idea is to change some subset H_0 of $[0, 1]$ -hospitals into $[1, 1]$ -hospitals so that in the optimal solution residents are assigned to hospitals

in H_0 plus original $[1,1]$ -hospitals exactly. However, there is no obvious way of selecting H_0 rather than exhaustive search, which will result in blow-ups of the computation time even if t is small. Furthermore, even if we would be able to find a correct H_0 , we are still not sure how to assign the residents to these (expanded) $[1,1]$ -hospitals optimally.

Theorem 3. *There is an $O(t^2(|H|(|R| + t))^{t+1})$ -time exact algorithm for 0-1 Min-BP HRMQ, where t is the cost of an optimal solution of a given instance.*

Proof. For a given integer $k > 0$, our algorithm $A(k)$ finds a solution (i.e., a matching between residents and hospitals) whose cost is at most k if any. Starting from $k = 1$, we run $A(k)$ until we find a solution, by increasing the value of k . $A(k)$ is quite simple, for which the following informal discussion might be enough.

Classify the blocking pairs as follows: A blocking pair (r, h) such that hospital h is empty is a Type-I blocking pair, and one such that h is non-empty is a Type-II blocking pair. First, guess the numbers k_1 of Type-I blocking pairs and k_2 of Type-II blocking pairs such that $k_1 + k_2 = k$, thus $k + 1$ possibilities.

We next guess, for each resident r_i , the number $b_i (\geq 0)$ of Type-I blocking pairs r_i is involved in, such that $b_1 + b_2 + \dots + b_{|R|} = k_1$. Note that there are $O((|R| + k_1)^{k_1})$ possibilities. Then, for each r_i such that $b_i > 0$, we again guess the set S_i of b_i hospitals that form Type-I blocking pairs with r_i . Note that there are at most $O(|H|^{k_1})$ different possibilities for selecting such k_1 blocking pairs. Let S be the union of S_i over all r_i . Note that all the hospitals in S are $[0, 1]$ -hospitals since they are supposed to be empty. Now remove these hospitals from the instance. For a resident r_i and hospital $h \in S$, suppose that $h \notin S_i$ in the current guess. Then, we remove all the hospitals lower ranked than h in r_i 's list, since if r_i is assigned to one of such hospitals, then r_i forms a Type-I blocking pair with h , contradicting our guess.

Next, we guess k_2 Type-II blocking pairs similarly. Let T be the set of those pairs. Since there are at most $|H||R|$ pairs, there are $O((|H||R|)^{k_2})$ choices of T . For each pair $(r, h) \in T$, we remove h from r 's list. Finally, we apply the Gale-Shapley algorithm to the modified instance. If all the $[1,1]$ -hospitals are full, then it is a desired solution, otherwise, we fail and proceed to the next guess.

We show that the algorithm runs correctly. Consider any optimal solution M_{opt} and consider the execution for $k = t$ for which we assume to have made a correct guess of the t blocking pairs of M_{opt} . Then, it is not hard to see that M_{opt} is stable in the above modified instance and makes all the $[1,1]$ -hospitals full. Then by the Rural Hospitals Theorem, any stable matching for this new instance makes all the $[1,1]$ -hospitals full. Hence if we apply the Gale-Shapley algorithm to this instance, we find a matching M in which all the $[1,1]$ -hospitals are full. Note that M has no blocking pair in the modified instance. Then, observe that M has at most k_1 Type-I blocking pairs in the original instance because, when hospitals in S are returned back to the instance, each resident r_i can form Type-I blocking pairs with only the hospitals in S_i . Also, only the pairs in T can be Type-II blocking pairs of M in the original instance. Therefore, M has at most t blocking pairs in the original instance.

Finally, we bound the time-complexity. For each k , we apply the Gale-Shapley algorithm to at most $O(k) \cdot O((|R| + k_1)^{k_1}) \cdot O(|H|^{k_1}) \cdot O((|H||R|)^{k_2}) = O(k(|H|(|R| + k))^{k_1})$ instances. Therefore, the time-complexity is $O(k(|H|(|R| + k))^{k_1})$ for each k . Since we find a solution when k is at most t , the whole time-complexity is at most $\sum_{k=1}^t O(k(|H|(|R| + k))^{k_1}) = O(t^2(|H|(|R| + t))^{t_1+1})$. $\square$

4 Minimum-Blocking-Resident HRMQ

In this section, we consider the problem of minimizing the number of blocking residents. We first show a negative result.

Theorem 4. *0-1 Min-BR 1ML-HRMQ is NP-hard. (Proof is omitted. See [14].)*

For the approximability, we note that Algorithm I in the proof of Theorem 2 does not work. For example, consider the instance introduced in Sec. 2. If we apply the Gale-Shapley algorithm, resident r_i is assigned to h_i for each i , and we need to move r_1 to h_{n+1} . However since h_1 becomes unassigned, all the residents become blocking residents. On the other hand, the optimal cost is 2 as we have seen there. Thus the approximation ratio becomes as bad as $\Omega(|R|)$.

Theorem 5. *There is a polynomial-time $\sqrt{|R|}$ -approximation algorithm for Min-BR HRMQ.*

We prove Theorem 5 by giving a $\sqrt{|R|}$ -approximation algorithm for 0-1 Min-BR HRMQ (Lemmas 1 and 3). In 0-1 Min-BR HRMQ, the number of residents assigned to each hospital is at most one. Hence, for a matching M , we sometimes abuse the notation $M(h)$ to denote the resident assigned to h (if any) although it is originally defined as the set of residents assigned to h .

Lemma 1. *If there is a polynomial-time α -approximation algorithm for 0-1 Min-BR HRMQ, then there is a polynomial-time α -approximation algorithm for Min-BR HRMQ. (Proof is omitted. See [14].)*

4.1 Our Algorithm

To describe the idea behind our algorithm, recall Algorithm I presented in the proof of Theorem 2: First, apply the Gale-Shapley algorithm to a given instance I and obtain a matching M . Next, move residents arbitrarily from assigned $[0, 1]$ -hospitals to unassigned $[1, 1]$ -hospitals. Suppose that in the course of the execution of Algorithm I, we move a resident r from a $[0, 1]$ -hospital h to an unassigned $[1, 1]$ -hospital. Then, of course r creates a blocking pair with h , but some other residents may also create blocking pairs with h because h becomes unassigned. Then, consider the following modification. First, set the upper bound quota of a $[0, 1]$ -hospital h to ∞ and apply the Gale-Shapley algorithm. Then, all residents who “wish” to go to h actually go there. Hence, even if we move all such residents to other hospitals, only the moved residents can become blocking residents. By doing this, we can bound the number of blocking residents by the

number (given by the function g introduced below) of those moving residents. In the above example, we extended the upper bound quota of only one hospital, but in fact, we may need to select two or more hospitals to select sufficiently many residents to be sent to other hospitals so as to make the matching feasible. However, at the same time, this number should be kept minimum to guarantee the quality of the solution.

We define $g(h, h)$: For an instance I of HR, suppose that we extend the upper bound quota of hospital h to ∞ and find a stable matching of this new instance. Define $g(h, h)$ be the number of residents assigned to h in this stable matching. Recall that this quantity does not depend on the choice of the stable matching by the Rural Hospitals Theorem [11]. Extend $g(h, h)$ to $g(A, B)$ for $A, B \subseteq H$ such that $g(A, B)$ denotes the number of residents assigned to hospitals in A when we change upper bound quotas of all the hospitals in B to ∞ .

We now propose Algorithm II for 0-1 Min-BR HRMQ. Let I be a given instance. Define $H_{p,q}$ to be the set of $[p, q]$ -hospitals of I . Recall from Sec. 3 that the deficiency of a hospital is the shortage of the assigned residents from its lower bound quota. Now define the *deficiency of the instance I* as the sum of the deficiencies of all the hospitals of I , and denote it $D(I)$. Since we are considering 0-1 Min-BR HRMQ, $D(I)$ is exactly the number of empty $[1, 1]$ -hospitals.

Algorithm II

- 1: Apply the Gale-Shapley algorithm to I by ignoring the lower bound quotas. Let M_s be the obtained matching. Compute the deficiency $D(I)$.
 - 2: $H'_{0,1} := \{h \mid M_s(h) \neq \emptyset, h \in H_{0,1}\}$. (If $M_s(h) = \emptyset$, then residents never come to h in this algorithm.)
 - 3: Compute $g(h, h)$ for each $h \in H'_{0,1}$ by using the Gale-Shapley algorithm.
 - 4: From $H'_{0,1}$, select $D(I)$ hospitals with smallest $g(h, h)$ values (ties are broken arbitrarily). Let S be the set of these hospitals. Extend the quota upper bounds of all hospitals in S to ∞ , and run the Gale-Shapley algorithm. Let M_∞ be the obtained matching.
 - 5: In M_∞ , move residents who are assigned to hospitals in S arbitrarily to empty hospitals to make the matching feasible. (We first make $[1, 1]$ -hospitals full. This is possible because of the restriction $\mathcal{Z}$. If there is a hospital in S still having two or more residents, then send remaining residents arbitrarily to empty $[0, 1]$ -hospitals, or simply make them unassigned if there is no $[0, 1]$ -hospital to send them.) Output the resulting matching M^* .
-

We first prove the following property of the original HR problem.

Lemma 2. *Let I_0 be an instance of HR, and h be any hospital. Let I_1 be a modification of I_0 so that only the upper bound quota of h is increased by 1. Let M_i be a stable matching of I_i for each $i \in \{0, 1\}$. Then, (i) $|M_0(h)| \leq |M_1(h)|$, and (ii) $\forall h' \in H \setminus \{h\}, |M_0(h')| \geq |M_1(h')|$.*

Proof. If M_0 is stable for I_1 , then we are done, so suppose not. Because M_0 is stable for I_0 , if M_0 has blocking pairs for I_1 , then all of them involve h . Let r be the resident such that (r, h) is a blocking pair and there is no blocking pair (r', h) such that h prefers r' to r . If we assign r to h , all blocking pairs including

h are removed. If no new blocking pairs arise, again, we are done. Otherwise, r must be previously assigned to some hospital, say h' , and all the new blocking pairs involve h' . We then choose the resident r' , most preferred by h' among all the blocking residents, and assign r' to h' . We continue this operation until there arise no new blocking pairs. This procedure eventually terminates because each iteration improves exactly one resident. By the termination condition, the resulting matching is stable for I_1 . Note that by this procedure, only h can gain one more resident, and at most one hospital may lose one resident. This completes the proof. $\square$

Obviously, Algorithm II runs in polynomial time. We show that Algorithm II runs correctly, namely that the output matching M^* satisfies the quota bounds. To do so, it suffices to show the following: (1) $|H'_{0,1}| \geq D(I)$ so that we can construct S at Step 4, and (2) in M_∞ , the total number of residents assigned to hospitals in S is at least the number of empty $[1, 1]$ -hospitals, so that Step 5 is executable.

For (1), recall that we have assumed that the number of residents is at least the sum of the lower bound quotas, which is the number of $[1, 1]$ -hospitals. Also, we can assume that all the residents are assigned in M_s (since otherwise, we already have a feasible stable matching). Then, by the definition of $D(I)$, we have that $|H'_{0,1}| \geq D(I)$. For (2), it suffices to show that the number N of residents assigned to $S \cup H_{1,1}$ in M_∞ is at least $|H_{1,1}|$. Note that empty hospitals in M_s are also empty in M_∞ by Lemma 2. Therefore, the number $\overline{N}$ of residents assigned to hospitals in $H \setminus (S \cup H_{1,1})$ in M_∞ is at most the number of hospitals in $H'_{0,1} \setminus S$. Thus $\overline{N} \leq |H'_{0,1}| - |S|$ and $N = |R| - \overline{N} \geq |R| - (|H'_{0,1}| - |S|)$. By the definition of $D(I)$, we have that $|H'_{0,1}| + |H_{1,1}| = |R| + D(I)$. Thus, $N \geq |R| - (|R| + D(I) - |H_{1,1}| - |S|) = |H_{1,1}|$ (recall that $|S| = D(I)$).

4.2 Analysis of the Approximation Ratio

Lemma 3. *The approximation ratio of Algorithm II is at most $\sqrt{|R|}$.*

Proof. Let I be a given instance of 0-1 Min-BR HRMQ and let f_{opt} and f_{alg} be the costs of an optimal solution and the solution obtained by Algorithm II, respectively. First, note that only residents moved at Step 5 can be blocking residents. Hence, (1) $f_{alg} \leq g(S, S)$.

We then give a lower bound on the optimal cost. To do so, see the proof of Theorem 2 [14], where it is shown that any optimal solution for instance I of Min-BP HRMQ has at least $D(I)$ blocking pairs. It should be noted that those $D(I)$ blocking pairs do not have any common resident. Thus we have (2) $f_{opt} \geq D(I)$.

Now here is our key lemma to evaluate the approximation ratio.

Lemma 4. *In Step 3 of Algorithm II, there are at least $D(I)$ different hospitals $h \in H'_{0,1}$ such that $g(h, h) \leq f_{opt}$.*

The proof will be given in a moment. By this lemma, we have $g(h, h) \leq f_{opt}$ for any $h \in S$ by the construction of S . This implies that (3) $\sum_{h \in S} g(h, h) \leq$

$D(I)f_{opt}$. Also, by Lemma 2, we have (4) $g(h, S) \leq g(h, h)$ for any $h \in S$. Hence, by (1), (4), (3) and (2), we have

$$f_{alg} \leq g(S, S) = \sum_{h \in S} g(h, S) \leq \sum_{h \in S} g(h, h) \leq D(I)f_{opt} \leq (f_{opt})^2.$$

Therefore, we have that $\sqrt{f_{alg}} \leq f_{opt}$, and hence $\frac{f_{alg}}{f_{opt}} \leq \sqrt{f_{alg}} \leq \sqrt{|R|}$, completing the proof of Lemma 3. $\square$

Proof of Lemma 4. Let h be a hospital satisfying the condition of the lemma. In order to calculate $g(h, h)$ in Step 3, we construct a stable matching, say M_h , for the instance $I_\infty(h)$ in which the upper bound quota of h is changed to ∞ . We do not know what kind of matching M_h is, but in the following, we show that there is a stable matching, say M_2 , for $I_\infty(h)$ such that $|M_2(h)| \leq f_{opt}$. M_h and M_2 may be different, but we can guarantee that $|M_h(h)| \leq f_{opt}$ by the Rural Hospitals Theorem. A bit trickily, we construct M_2 from an optimal matching.

Let M_{opt} be an optimal solution of I (which of course we do not know). Let R_b and R_n be the sets of blocking residents and non-blocking residents for M_{opt} , respectively. Then $|R_b| = f_{opt}$ by definition. We modify M_{opt} as follows: Take any resident $r \in R_b$. If r is unassigned, we do nothing. Otherwise, force r to be unassigned. Then there may arise new blocking pairs, all of which include the hospital h' to which r was assigned. Among residents who are included in such new blocking pairs, we select the resident in R_n who is most preferred by h' (if any) and assign her to h' . In a similar way as the proof of Lemma 2, we continue to move residents until no new blocking pair arises (but this time, we move only residents in R_n as explained above). We do this for all the residents in R_b , and let M_1 be the resulting matching.

The following (a) and (b) are immediate: (a) There are at least f_{opt} unassigned residents in M_1 , since residents in R_b are unassigned in M_1 . (b) Residents in R_n are non-blocking for M_1 . We prove the following properties: (c) There are at most f_{opt} unassigned $[1, 1]$ -hospitals in M_1 . (d) Define $H' = \{h \mid h \in H'_{0,1} \text{ and } h \text{ is unassigned in } M_1\}$. Then $|H'| \geq D(I)$.

(c) In M_{opt} , all the $[1, 1]$ -hospitals are full. It is easy to see that an unassigned hospital of M_{opt} is also unassigned in M_1 . Since at most f_{opt} residents are made to be unassigned by the above procedure, the claim holds.

(d) Let H_1 be the set of hospitals assigned in M_1 . By the definition of H' , $H' = H'_{0,1} \setminus (H_1 \cap H_{0,1})$. By the definition of $D(I)$, $|H'_{0,1}| = |R| + D(I) - |H_{1,1}|$, and by the properties (a) and (c), $|H_1 \cap H_{0,1}| \leq |R| - |H_{1,1}|$. Then $|H'| \geq |H'_{0,1}| - |H_1 \cap H_{0,1}| \geq (|R| + D(I) - |H_{1,1}|) - (|R| - |H_{1,1}|) = D(I)$.

For any $h \in H'$, we show that $g(h, h) \leq f_{opt}$. Then, this completes the proof of Lemma 4 because $H' \subseteq H'_{0,1}$ and (d) $|H'| \geq D(I)$. Since h is unassigned in M_1 , residents in R_n are still non-blocking for M_1 in $I_\infty(h)$ (whose definition is in the beginning of this proof) by the property (b). Now, choose any resident r from R_b , and apply the Gale-Shapley algorithm to $I_\infty(h)$ starting from M_1 . This execution starts from the proposal by r , and at the end, nobody in $R_n \cup \{r\}$ is a blocking resident for $I_\infty(h)$. Since hospitals assigned in M_1 never become unassigned, and

since unassigned residents in R_n never become assigned, h receives at most one resident. If we do this for all the residents in R_b , the resulting matching M_2 is stable for $I_\infty(h)$, and h is assigned at most $|R_b| = f_{opt}$ residents. As mentioned previously, this implies $g(h, h) \leq f_{opt}$. $\square$

We can show that the analysis of Lemma 3 is tight up to a constant factor: There is an instance of 0-1 Min-BR HRMQ for which Algorithm II produces a solution of cost $|R| - \sqrt{|R|}$ but the optimal cost is at most $2\sqrt{|R|}$ (see [14]).

4.3 Inapproximability of Min-BR HRMQ

For the hardness of Min-BR HRMQ, we have only NP-hardness, but we can give a strong evidence for its inapproximability. The *Dense k -Subgraph Problem* (DkS) is the problem of finding, given a graph G and a positive integer k , an induced subgraph of G with k vertices that contains as many edges as possible. This problem is NP-hard because it is a generalization of Max CLIQUE. Its approximability has been studied intensively but there still remains a large gap between approximability and inapproximability: The best known approximation ratio is $|V|^{1/4+\epsilon}$ [4], while there is no PTAS under reasonable assumptions [8, 19]. The following Theorem 6 shows that approximating Min-BR HRMQ within a constant ratio implies the same for DkS.

Theorem 6. *If Min-BR 1ML-HRMQ has a polynomial-time c -approximation algorithm, then DkS has a polynomial-time $(1 + \epsilon)c^4$ -approximation algorithm for any positive constant ϵ . (Proof is omitted. See [14].)*

5 Concluding Remarks

An obvious future research is to obtain lower bounds on the approximation factor for Min-BR HRMQ (we even do not know its APX-hardness at this moment). Since the problem is harder than DkS, it should be a reasonable challenge. Another direction is to develop an FPT algorithm for Min-BP HRMQ, improving Theorem 3. Finally, we remark on the possibility of generalization of instances: In this paper, we guarantee existence of feasible matchings by the restriction $\mathcal{Z}$ (Sec. 2). However, even if we allow arbitrarily incomplete lists (and even ties), it is decidable in polynomial time if the given instance admits a feasible matching [9]. Thus, it might be interesting to seek approximate solution for instances without restriction $\mathcal{Z}$. Unfortunately, however, we can easily imply its $|R|^{1-\epsilon}$ -hardness (see [14]).

References

1. D. J. Abraham, P. Biró and D. F. Manlove, “ “Almost stable” matchings in the roommates problem,” *Proc. WAOA 2005*, LNCS 3879, pp. 1–14, 2005.
2. D. J. Abraham, R. W. Irving and D. F. Manlove, “Two algorithms for the Student-Project Allocation problem,” *J. Discrete Algorithms*, Vol.5(1), pp. 73–90, 2007.

3. B. Aldershof and O. M. Carducci, "Stable matchings with couples," *Discrete Applied Mathematics*, Vol. 68, pp. 203–207, 1996.
4. A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan, "Detecting high log-densities – an $O(n^{1/4})$ approximation for densest k -subgraph," *Proc. STOC 2010*, pp. 201–210, 2010.
5. P. Biró, T. Fleiner, R. W. Irving, and D. F. Manlove, "The College Admissions problem with lower and common quotas," *Theoretical Computer Science*, Vol. 411(34-36), pp. 3136–3153, 2010.
6. P. Biró, D. F. Manlove, and S. Mittal, "Size versus stability in the marriage problem," *Theoretical Computer Science*, Vol. 411(16-18), pp. 1828–1841, 2010.
7. Canadian Resident Matching Service (CaRMS), <http://www.carms.ca/>
8. U. Feige, "Relations between average case complexity and approximation complexity," In *Proc. STOC 2002*, pp. 534–543, 2002.
9. H. N. Gabow, "An efficient reduction technique for degree-constrained subgraph and bidirected network flow problems," In *Proc. STOC 83*, pp. 448–456, 1983.
10. D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *Amer. Math. Monthly*, Vol. 69, pp. 9–15, 1962.
11. D. Gale and M. Sotomayor, "Some remarks on the stable matching problem," *Discrete Applied Mathematics*, Vol. 11, pp. 223–232, 1985.
12. D. Gusfield and R. W. Irving, "The Stable Marriage Problem: Structure and Algorithms," MIT Press, Boston, MA, 1989.
13. K. Hamada, K. Iwama, and S. Miyazaki, "The hospitals/residents problem with quota lower bounds," *Proc. MATCH-UP* (satellite workshop of ICALP 2008), pp. 55–66, 2008.
14. K. Hamada, K. Iwama, and S. Miyazaki, "The hospitals/residents problem with quota lower bounds," manuscript. <http://www.lab2.kuis.kyoto-u.ac.jp/~shuichi/ESA11/esa11-final-long.pdf>
15. K. Hamada, K. Iwama, and S. Miyazaki, "An improved approximation lower bound for finding almost stable maximum matchings," *Information Processing Letters*, Vol. 109(18), pp. 1036–1040, 2009.
16. C.-C. Huang, "Classified stable matching," *Proc. SODA 2010*, pp. 1235–1253, 2010.
17. R. W. Irving, D. F. Manlove and S. Scott, "The hospital/residents problem with ties," *Proc. SWAT 2000*, LNCS 1851, pp. 259–271, 2000.
18. R. W. Irving, D. F. Manlove and S. Scott, "The stable marriage problem with master preference lists," *Discrete Applied Math.*, Vol. 156(15), pp. 2959–2977, 2008.
19. S. Khot, "Ruling out PTAS for graph min-bisection, densest subgraph and bipartite clique," In *Proc. FOCS 2004*, pp. 136–145, 2004.
20. S. Khuller, S. G. Mitchell and V. V. Vazirani, "On-Line algorithms for weighted bipartite matching and stable marriages," *Theoretical Computer Science*, Vol. 127(2), pp. 255–267, 1994.
21. U. Feige, G. Kortsarz, and D. Peleg, "The dense k -subgraph problem," *Algorithmica*, Vol. 29, pp. 410–421, 2001.
22. E. J. McDermid and D. F. Manlove, "Keeping partners together: algorithmic results for the hospitals/residents problem with couples," *Journal of Combinatorial Optimization*, Vol. 19(3), pp. 279–303, 2010.
23. E. Ronn, "NP-complete stable matching problems," *J. Algorithms*, Vol. 11, pp. 285–304, 1990.
24. A. E. Roth, "The evolution of the labor market for medical interns and residents: a case study in game theory," *J. Political Economy*, Vol. 92(6), pp. 991–1016, 1984.
25. S. A. Vinterbo, "A stab at approximating minimum subadditive join," *Proc. WADS 2007*, LNCS 4619, pp. 214–225, 2007.