



University
of Glasgow

Manlove, D.F. *Hospitals / residents problem*. In Kao, M-Y. (Eds)
Encyclopedia of Algorithms, Chap 180, pages pp. 390-394. New York:
Springer (2008)

<http://eprints.gla.ac.uk/4495/>

Deposited on: 18 July 2008

The Hospitals / Residents Problem (1962; Gale, Shapley)

David F. Manlove, University of Glasgow, www.dcs.gla.ac.uk/~davidm

entry editor: Kazuo Iwama

SYNONYMS: College Admissions problem; University Admissions problem; Stable Admissions problem; Stable Assignment problem; Stable b -matching problem

INDEX TERMS: Matching; stability

1 PROBLEM DEFINITION

An instance I of the Hospitals / Residents problem (HR) [5, 6, 14] involves a set $R = \{r_1, \dots, r_n\}$ of *residents* and a set $H = \{h_1, \dots, h_m\}$ of *hospitals*. Each hospital $h_j \in H$ has a positive integral *capacity*, denoted by c_j . Also, each resident $r_i \in R$ has a *preference list* in which he ranks in strict order a subset of H . A pair $(r_i, h_j) \in R \times H$ is said to be *acceptable* if h_j appears in r_i 's preference list; in this case r_i is said to *find* h_j *acceptable*. Similarly each hospital $h_j \in H$ has a preference list in which it ranks in strict order those residents who find h_j acceptable. Given any three agents $x, y, z \in R \cup H$, x is said to *prefer* y to z if x finds each of y and z acceptable, and y precedes z on x 's preference list. Let $C = \sum_{h_j \in H} c_j$.

Let A denote the set of acceptable pairs in I , and let $L = |A|$. An *assignment* M is a subset of A . If $(r_i, h_j) \in M$, r_i is said to be *assigned to* h_j , and h_j is *assigned* r_i . For each $q \in R \cup H$, the set of assignees of q in M is denoted by $M(q)$. If $r_i \in R$ and $M(r_i) = \emptyset$, r_i is said to be *unassigned*, otherwise r_i is *assigned*. Similarly, any hospital $h_j \in H$ is *under-subscribed*, *full* or *over-subscribed* according as $|M(h_j)|$ is less than, equal to, or greater than c_j , respectively.

A *matching* M is an assignment such that $|M(r_i)| \leq 1$ for each $r_i \in R$ and $|M(h_j)| \leq c_j$ for each $h_j \in H$ (i.e., no resident is assigned to an unacceptable hospital, each resident is assigned to at most one hospital, and no hospital is over-subscribed). For notational convenience, given a matching M and a resident $r_i \in R$ such that $M(r_i) \neq \emptyset$, where there is no ambiguity the notation $M(r_i)$ is also used to refer to the single member of $M(r_i)$.

A pair $(r_i, h_j) \in A \setminus M$ *blocks* a matching M , or is a *blocking pair* for M , if the following conditions are satisfied relative to M :

1. r_i is unassigned or prefers h_j to $M(r_i)$;
2. h_j is under-subscribed or prefers r_i to at least one member of $M(h_j)$ (or both).

A matching M is said to be *stable* if it admits no blocking pair. Given an instance I of HR, the problem is to find a stable matching in I .

2 KEY RESULTS

HR was first defined by Gale and Shapley [5] under the name “College Admissions Problem”. In their seminal paper, the authors’ primary consideration is the classical Stable

```

 $M := \emptyset;$ 
while (some resident  $r_i$  is unassigned and  $r_i$  has a non-empty list) {
     $h_j :=$  first hospital on  $r_i$ 's list;
    /*  $r_i$  applies to  $h_j$  */
     $M := M \cup \{(r_i, h_j)\};$ 
    if ( $h_j$  is over-subscribed) {
         $r_k :=$  worst resident in  $M(h_j)$  according to  $h_j$ 's list;
         $M := M \setminus \{(r_k, h_j)\};$ 
    }
    if ( $h_j$  is full) {
         $r_k :=$  worst resident in  $M(h_j)$  according to  $h_j$ 's list;
        for (each successor  $r_l$  of  $r_k$  on  $h_j$ 's list)
            delete the pair  $(r_l, h_j);$ 
    }
}

```

Figure 1: Gale/Shapley algorithm for HR.

Marriage problem (SM; see Entries 142 and 143), which is a special case of HR in which $n = m$, $A = R \times H$, and $c_j = 1$ for all $h_j \in H$ – in this case, the residents and hospitals are more commonly referred to as the men and women respectively. Gale and Shapley show that every instance I of HR admits at least one stable matching. Their proof of this result is constructive, i.e., an algorithm for finding a stable matching in I is described. This algorithm has become known as the *Gale/Shapley algorithm*.

An extended version of the Gale/Shapley algorithm for HR is shown in Figure 1. The algorithm involves a sequence of *apply* and *delete* operations. At each iteration of the while loop, some unassigned resident r_i with a non-empty preference list applies to the first hospital h_j on his list, and becomes provisionally assigned to h_j (this assignment could subsequently be broken). If h_j becomes over-subscribed as a result of this assignment, then h_j rejects its worst assigned resident r_k . Next, if h_j is full (irrespective of whether h_j was over-subscribed earlier in the same loop iteration), then for each resident r_l that h_j finds less desirable than its worst resident r_k , the algorithm *deletes the pair* (r_l, h_j) , which comprises deleting h_j from r_l 's preference list and vice versa.

Given that the above algorithm involves residents applying to hospitals, it has become known as the *Resident-oriented* Gale/Shapley algorithm, or RGS algorithm for short [6, Section 1.6.3]. The RGS algorithm terminates with a stable matching, given an instance of HR [5] [6, Theorem 1.6.2]. Using a suitable choice of data structures (extending those described in [6, Section 1.2.3]), the RGS algorithm can be implemented to run in $O(L)$ time. This algorithm produces the stable matching that is simultaneously best-possible for all residents [5] [6, Theorem 1.6.2]. These observations may be summarised as follows:

Theorem 2.1. *Given an instance of HR, the RGS algorithm constructs, in $O(L)$ time, the unique stable matching in which each assigned resident obtains the best hospital that he could obtain in any stable matching, whilst each unassigned resident is unassigned in every stable matching.*

A counterpart of the RGS algorithm, known as the *Hospital-oriented* Gale/Shapley algorithm, or HGS algorithm for short [6, Section 1.6.2], gives the unique stable matching that similarly satisfies an optimality property for the hospitals [6, Theorem 1.6.1].

Although there may be many stable matchings for a given instance I of HR, some key structural properties hold regarding unassigned residents and under-subscribed hospitals with respect to all stable matchings in I , as follows.

Theorem 2.2. *For a given instance of HR,*

- *the same residents are assigned in all stable matchings;*
- *each hospital is assigned the same number of residents in all stable matchings;*
- *any hospital that is under-subscribed in one stable matching is assigned exactly the same set of residents in all stable matchings.*

These results are collectively known as the “Rural Hospitals Theorem” (see [6, Section 1.6.4] for further details). Furthermore, the set of stable matchings in I forms a distributive lattice under a natural dominance relation [6, Section 1.6.5].

3 APPLICATIONS

Practical applications of HR are widespread, most notably arising in the context of centralised automated matching schemes that assign applicants to posts (for example medical students to hospitals, school-leavers to universities, and primary school pupils to secondary schools). Perhaps the largest and best-known example of such a scheme is the National Resident Matching Program (NRMP) in the US [16], which annually assigns around 31,000 graduating medical students (known as residents) to their first hospital posts, taking into account the preferences of residents over hospitals and vice versa, and the hospital capacities. Counterparts of the NRMP are in existence in other countries, including Canada [17], Scotland [18] and Japan [19]. These matching schemes essentially employ extensions of the RGS algorithm for HR.

Centralised matching schemes based largely on HR also occur in other practical contexts, such as school placement in New York [1], university faculty recruitment in France [3] and university admission in Spain [12].

Extensions of HR. One key extension of HR that has considerable practical importance arises when an instance may involve a set of couples, each of which submits a joint preference list over pairs of hospitals (typically in order that the members of a given couple can be located geographically close to one another, for example). The extension of HR in which couples may be involved is denoted by HRC; the stability definition in HRC is a natural extension of that in HR (see [6, Section 1.6.6] for a formal definition of HRC). It is known that an instance of HRC need not admit a stable matching (see [6, Section 1.6.6] and [14, Section 5.4.3]). Moreover, the problem of deciding whether an HRC instance admits a stable matching is NP-complete [13].

HR may be regarded as a many-one generalisation of SM. A further generalisation of SM is to a many-many stable matching problem, in which both residents and hospitals may be multiply assigned subject to capacity constraints. In this case, residents and hospitals are more commonly referred to as workers and firms respectively. There are two basic variations of the many-many stable matching problem according to whether (i) workers rank acceptable firms in order of preference and vice versa, or (ii) workers rank acceptable *subsets* of firms in order of preference and vice versa. Previous work relating to both models is surveyed in [4].

Other variants of HR may be obtained if preference lists include ties. This extension is again important from a practical perspective, since it may be unrealistic to expect a popular hospital to rank a large number of applicants in strict order, particularly if it is indifferent among groups of applicants. The extension of HR in which preference lists may include ties is denoted by HRT. In this context three natural stability definitions arise, so-called *weak stability*, *strong stability* and *super-stability* (see [8] for formal definitions of these concepts). Given an instance I of HRT, it is known that weakly stable matchings may have different sizes, and the problem of finding a maximum cardinality weakly stable

matching is NP-hard (see Entry 151 for further details). On the other hand, in contrast to the case for weak stability, a super-stable matching in I need not exist, though there is an $O(L)$ algorithm to find a such a matching if one does [7]. Analogous results hold in the case of strong stability – in this case an $O(L^2)$ algorithm [8] was improved by an $O(CL)$ algorithm [10] and extended to the many-many case [11]. Furthermore, counterparts of the Rural Hospitals Theorem hold for HRT under each of the super-stability and strong stability criteria [7, 15].

A further generalisation of HR arises when each hospital may be split into several departments, where each department has a capacity, and residents rank individual departments in order of preference. This variant is modelled by the Student-Project Allocation problem [2]. Finally, the Stable Fixtures problem [9] is a non-bipartite extension of HR in which there is a single set of agents, each of whom has a capacity and ranks a subset of the others in order of preference.

4 OPEN PROBLEMS

Several approximation algorithms for finding a maximum cardinality weakly stable matching have been formulated, given an instance of HRT where each hospital has capacity 1 (see Entry 151 for further details). It remains open to extend these algorithms or to formulate effective heuristics for the case of HRT with arbitrary capacities. This problem is particularly relevant from the practical perspective, since as already noted in Section 3, hospitals may wish to include ties in their preference lists. In this case weak stability is the most commonly-used stability criterion, due to the guaranteed existence of such a matching. Attempting to match as many residents as possible motivates the search for large weakly stable matchings.

5 EXPERIMENTAL RESULTS

None is reported.

6 DATA SETS

None is reported.

7 URL to CODE

Ada implementations of the RGS and HGS algorithms for HR may be found via the following URL: <http://www.dcs.gla.ac.uk/research/algorithms/stable>.

8 CROSS REFERENCES

See also Entries 142, 143, 151, 153, 154, 163 and 436 of this volume.

9 RECOMMENDED READING

- [1] A. Abdulkadiroğlu, P.A. Pathak, and A.E. Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, 2006.
- [2] D.J. Abraham, R.W. Irving, and D.F. Manlove. Two algorithms for the Student-Project allocation problem. *Journal of Discrete Algorithms*, 5(1):73–90, 2007.
- [3] M. Baïou and M. Balinski. Student admissions and faculty recruitment. *Theoretical Computer Science*, 322(2):245–265, 2004.

- [4] V. Bansal, A. Agrawal, and V.S. Malhotra. Stable marriages with multiple partners: efficient search for an optimal solution. In *Proceedings of ICALP '03: the 30th International Colloquium on Automata, Languages and Programming*, volume 2719 of *Lecture Notes in Computer Science*, pages 527–542. Springer, 2003.
- [5] D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [6] D. Gusfield and R.W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.
- [7] R.W. Irving, D.F. Manlove, and S. Scott. The Hospitals/Residents problem with Ties. In *Proceedings of SWAT 2000: the 7th Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 259–271. Springer, 2000.
- [8] R.W. Irving, D.F. Manlove, and S. Scott. Strong stability in the Hospitals/Residents problem. In *Proceedings of STACS 2003: the 20th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2607 of *Lecture Notes in Computer Science*, pages 439–450. Springer, 2003.
- [9] R.W. Irving and S. Scott. The stable fixtures problem. To appear in *Discrete Applied Mathematics*, 2007.
- [10] T. Kavitha, K. Mehlhorn, D. Michail, and K. Paluch. Strongly stable matchings in time $O(nm)$ and extension to the Hospitals-Residents problem. In *Proceedings of STACS 2004: the 21st International Symposium on Theoretical Aspects of Computer Science*, volume 2996 of *Lecture Notes in Computer Science*, pages 222–233. Springer, 2004.
- [11] V.S. Malhotra. On the stability of multiple partner stable marriages with ties. In *Proceedings of ESA '04: the Twelfth Annual European Symposium on Algorithms*, volume 3221 of *Lecture Notes in Computer Science*, pages 508–519. Springer, 2004.
- [12] A. Romero-Medina. Implementation of stable solutions in a restricted matching market. *Review of Economic Design*, 3(2):137–147, 1998.
- [13] E. Ronn. NP-complete stable matching problems. *Journal of Algorithms*, 11:285–304, 1990.
- [14] A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis*, volume 18 of *Econometric Society Monographs*. Cambridge University Press, 1990.
- [15] S. Scott. *A study of stable marriage problems with ties*. PhD thesis, University of Glasgow, Department of Computing Science, 2005.
- [16] <http://www.nrmp.org/> (National Resident Matching Program website).
- [17] <http://www.carms.ca> (Canadian Resident Matching Service website).
- [18] <http://www.nes.scot.nhs.uk/sfas/> (Scottish Foundation Allocation Scheme website).
- [19] <http://www.jrmp.jp> (Japan Resident Matching Program website).