
MŰHELYTANULMÁNYOK

DISCUSSION PAPERS

MT-DP – 2013/30

Matching Couples with Scarf's Algorithm

PÉTER BIRÓ - TAMÁS FLEINER - ROB IRVING

Discussion papers
MT-DP – 2013/30

Institute of Economics, Centre for Economic and Regional Studies,
Hungarian Academy of Sciences

KTI/IE Discussion Papers are circulated to promote discussion and provoke comments.
Any references to discussion papers should clearly state that the paper is preliminary.
Materials published in this series may subject to further publication.

Matching Couples with Scarf's Algorithm

Authors:

Péter Biró
research fellow
Institute of Economics
Centre for Economic and Regional Studies
Hungarian Academy of Sciences
email: biro.peter@krtk.mta.hu

Tamás Fleiner
associate professor
Dept. of C.S. and Inf. Theory
Budapest Univ. of Technology and Economics
email: fleiner@cs.bme.hu

Rob Irving
honorary research fellow
School of Computing Science
University of Glasgow
Rob.Irving@glasgow.ac.uk

September 2013

ISBN 978-615-5243-89-9
ISSN 1785 377X

Matching Couples with Scarf's Algorithm

Péter Biró - Tamás Fleiner - Rob Irving

Abstract

Scarf's algorithm [18] provides fractional core elements for NTU-games. Biró and Fleiner [3] showed that Scarf's algorithm can be extended for capacitated NTU-games. In this setting agents can be involved in more than one coalition at a time, cooperations may be performed with different intensities up to some limits, and the contribution of the agents can also differ in a coalition. The fractional stable solutions for the above model, produced by the extended Scarf algorithm, are called stable allocations. In this paper we apply this solution concept for the Hospitals Residents problem with Couples (HRC). This is one of the most important general stable matching problems due to its relevant applications, also well-known to be NP-hard. We show that if a stable allocation yielded by the Scarf algorithm turns out to be integral then it provides a stable matching for an instance of HRC, so this method can be used as a heuristic. In an experimental study, we compare this method with other heuristics constructed for HRC that are applied in practice in the American and Scottish resident allocation programs, respectively. Our main finding is that the Scarf algorithm outperforms all the other known heuristics when the proportion of couples is high.

Keywords: Scarf lemma, stable allocation, hospitals residents problem, couples

JEL classification: C71, C78

Acknowledgement:

Supported by OTKA grant no. K108383. The author is a member of the MTA-ELTE Egervary Research Group. The research was supported by grant (no. CK 80124) from the National Development Agency of Hungary, based on a source from the Research and Technology Innovation Fund and the results discussed below are supported by the grant TÁMOP - 4.2.2.B-10/1-2010-0009.

Házastársak párosítása Scarf algoritmusával

Biró Péter - Fleiner Tamás - Rob Irving

Összefoglaló

Scarf algoritmus a tört-mag elemeket ad eredményül NTU-játékokra. Biró és Fleiner megmutatta, hogy Scarf algoritmus kiterjeszthető kapacitásos NTU-játékokra. Ebben a modellben a játékosok részt vehetnek egynél több koalícióban, az együttműködések volumene változtatható egy adott értékig, illetve a játékosok részesedése is különbözhet egy koalícióban. A fenti problémára mindig létező tört-stabil megoldásokat stabil allokációnak hívjuk, ez kapható meg a kiterjesztett Scarf algoritmussal. Ebben a tanulmányban e modellt és megoldási koncepciót alkalmazzuk a rezidensfelvételi probléma házastársas változatára. A gyakorlati alkalmazások miatt ez az egyik legfontosabb stabil párosítási probléma, amely viszont bizonyítottan NP-nehéz. Megmutatjuk, hogy ha a kiterjesztett Scarf algoritmussal kapott megoldás egészértékű, akkor megfelel egy stabil megoldásnak, vagyis ez az eljárás alkalmazható heurisztikaként. Egy szimulációs esettanulmányban összehasonlítjuk ezt a módszert más heurisztikákkal, amelyek az amerikai és a skóciai alkalmazásokban működnek. A legfontosabb megfigyelésünk, hogy a Scarf algoritmuson alapuló heurisztika jelentősen felülmúlja a többit abban az esetben, amikor a házaspárok aránya a mintában magas.

Tárgyszavak: Scarf lemma, stabil allokáció, rezidens allokációs probléma, házastársak

JEL kódok: C71, C78

Matching Couples with Scarf's Algorithm

PÉTER BIRÓ¹

Institute of Economics, RCERS-HAS
and
Dept. of Op. Res. and Act. Sci., CUB
1112. Budapest, Budaörsi út 45., Hungary
peter.biro@krtk.mta.hu

TAMÁS FLEINER²

Dept. of C.S. and Inf. Theory
Budapest Univ. of Technology and Economics
1117 Budapest, Magyar tudósok körútja 2.,
Hungary
fleiner@cs.bme.hu

ROB IRVING

School of Computing Science
University of Glasgow
Sir Alwyn Williams Building,
Glasgow G12 8QQ, UK
Rob.Irving@glasgow.ac.uk

Abstract: Scarf's algorithm [18] provides fractional core elements for NTU-games. Biró and Fleiner [3] showed that Scarf's algorithm can be extended for capacitated NTU-games. In this setting agents can be involved in more than one coalition at a time, cooperations may be performed with different intensities up to some limits, and the contribution of the agents can also differ in a coalition. The fractional stable solutions for the above model, produced by the extended Scarf algorithm, are called stable allocations. In this paper we apply this solution concept for the Hospitals Residents problem with Couples (HRC). This is one of the most important general stable matching problems due to its relevant applications, also well-known to be NP-hard. We show that if a stable allocation yielded by the Scarf algorithm turns out to be integral then it provides a stable matching for an instance of HRC, so this method can be used as a heuristic. In an experimental study, we compare this method with other heuristics constructed for HRC that are applied in practice in the American and Scottish resident allocation programs, respectively. Our main finding is that the Scarf algorithm outperforms all the other known heuristics when the proportion of couples is high.

Keywords: Scarf lemma, stable allocation, hospitals residents problem, couples

JEL codes: C71, C78

1 Introduction

Mechanism design in matching markets dates back to the seminal paper of Gale and Shapley [7] on college admissions. They introduced the concept of *stable matching*, that is a fair solution where an application of a student can be rejected by a college only if its quota is filled with better candidates. Gale and Shapley gave an efficient algorithm to find a stable matching in this setting. It turned out [15]

¹Corresponding author. Research is supported by the Hungarian Academy of Sciences under its Momentum Programme (LD-004/2010) and also by OTKA grant no. K108673.

²Supported by OTKA grant no. K108383. The author is a member of the MTA-ELTE Egerváry Research Group. The research was supported by grant (no. CK 80124) from the National Development Agency of Hungary, based on a source from the Research and Technology Innovation Fund and the results discussed below are supported by the grant TÁMOP - 4.2.2.B-10/1-2010-0009.

that the same method had already been used in the US resident matching program (NRMP) since 1952. This program has been redesigned later [16], partly because the organisers wanted to accommodate the wishes of couples. Since then couples can submit joint preference lists in order to avoid being matched to hospitals far from each other. However, for the *Hospitals Residents problem with couples* (HRC) the existence of a stable matching is no longer guaranteed [15]. Moreover, the related decision problem is NP-hard [14], therefore we need to use heuristics for large markets. There are many recent papers on this problem by economists [16], [9], [10], [11], and by computer scientists [1], [19], [12], [13] as well, see also an interdisciplinary survey [5]. In particular, Biró, Irving and Schlotter [4] compared some old and new heuristics for a setting that is currently present in the Scottish resident allocation program [21, 23].

Another seminal paper in cooperative game theory is by Scarf [18]. He gave an algorithm to find a core element for any *balanced NTU-game*. Aharoni and Fleiner [2] used this algorithm to find *stable fractional matchings* for problems where the underlying graph is not necessarily bipartite. Biró and Fleiner [3] generalised this result by showing that Scarf’s algorithm can be extended to find *stable allocations* for NTU-games where both the agents and their cooperations can have capacities, and even when agents have different contributions in a co-operation performed. We shall note though that the Scarf algorithm is not known to be polynomial.

In this paper first we show that the HRC problem where couples may not apply for a pair of positions in the same hospital, can be reduced to an integral stable allocation problem (ISA) with no edge capacities, where we have a one-to-one correspondence between the stable solutions of the two models. Furthermore, we show that the HRC problem (where couples may apply for a pair of positions in the same hospital) can be transformed into an integral stable allocation problem with contributions (ISAC) and with no edge capacities, where the stable solutions for the latter problem are stable solutions for the former. Biró and Fleiner [3] demonstrated that the original Scarf algorithm always returns a stable allocation for both the ISA and ISAC problems with no edge capacities, and whenever it outputs an integer solution it solves the problem. Therefore, the Scarf algorithm can be used as a heuristic to solve all of these problems, ISA, ISAC and HRC as well.

After this, we present an experimental study, that follows up the work by Biró, Irving and Schlotter [4], where we compare the performance of the Scarf algorithm and other heuristics described in [16] and [4]. Our main finding is that the Scarf algorithm works very well if the proportion of couples is high. However, it is unrealistic to suppose that all the applicants form couples in a resident allocation program, but there can be other applications with a similar feature. For instance, in the Hungarian higher education matching scheme [20] students can apply to pairs of teacher courses, and they usually do so in a very concentrated way. This causes a very similar problem, as a student applying to pairs of courses can be seen as a couple applying to pairs of positions. The same situation occurred in the Scottish resident allocation scheme (SPA) from 2000 to 2005, where the medical doctors had to apply for two posts, a medical and a surgical one [8], although the elicitation of their preferences was more restricted than in the Hungarian application.

2 Description of the models

First we present a model for the hospitals residents problem with couples, similar to the one described by Biró, Irving and Schlotter [4]. Then we present the integral stable allocation problem with contributions by Biró and Fleiner [3]. Finally we describe the connections between some variants of these two problems.

2.1 The Hospitals Residents problem with Couples (HRC)

In the Hospitals Residents problem with Couples (HRC) we are given a set of applicants $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$ and a set of hospitals $\mathcal{H} = \{h_1, h_2 \dots h_m\}$ with $c(h_p)$ denoting the capacity of hospital h_p . The set of applicants can be partitioned into single applicants \mathcal{S} and couples \mathcal{C} . The preference list of a single applicant $a_i \in \mathcal{S}$ contains hospitals in the order of her preferences, whilst the preference list of couple $(a_i, a_j) \in \mathcal{C}$ contains pairs of hospitals in the order of their joint preferences. An application of a single

applicant to a hospital is referred to as a *single application*, a *joint application* is made by a couple to a pair of hospitals. In our general model we allow that a couple may apply for a pair of positions in the same hospital, in which case we refer to this application as a *combined application*. Let us denote the above three types of applications, single applications, joint applications, and combined applications by E^S , E^J and E^C , respectively. For simplicity we assume that the residents apply to hospitals that are acceptable for them and every resident applying to a hospital is acceptable for that hospital. The solution is a matching M that consists of employments of form $\langle a_i, h_p \rangle$, where if $a_i \in \mathcal{S}$ then $[a_i \rightarrow h_p] \in E^S$, and if $(a_i, a_j) \in \mathcal{C}$ then either $[(a_i, a_j) \rightarrow (h_p, h_q)] \in E^J$ and also $\langle a_j, h_q \rangle \in M$, or $[(a_i, a_j) \rightarrow (h_p, h_p)] \in E^I$ and also $\langle a_j, h_p \rangle \in M$. Let $M(a_i)$ denote the hospital where a_i is allocated in M , if any, and let $M(h_p)$ denote the set of applicants allocated to h_p . A matching has to respect the capacity constraints, i.e. no applicant may be allocated to more than one hospital and no hospital may employ more residents than its capacity, so $|M(h_p)| \leq c(h_p)$. We say that h_p is *full* with respect to M if $|M(h_p)| = c(h_p)$ and *undersubscribed* otherwise. Finally we note that we can also accommodate the possibility that only one member of a couple applies for a position (and the other member remains unmatched) in our model by simply introducing a dummy hospital with no capacity constraint, which would correspond to the outside option.

To define stability we need to specify the preferences of the hospitals over the set of applications. We suppose that every hospital has a strict preference ordering over the acceptable applicants. From these rankings we will derive the definition of stability, that is a slight extension of the stability definition by Biró, Irving and Schlotter [4]¹. A matching M is *stable* if it is not *blocked* by a pair $\langle a_i, h_p \rangle$ consisting of a single applicant a_i and a hospital h_p , or by a coalition $\langle (a_i, a_j), (h_p, h_q) \rangle$ consisting of a couple (a_i, a_j) and distinct hospitals h_p and h_q , or by a coalition $\langle (a_i, a_j), (h_p, h_p) \rangle$ consisting of a couple (a_i, a_j) and a hospital h_p .

A single applicant a_i and a hospital h_p *block* M if

- (a) a_i is unmatched, or prefers h_p to $M(a_i)$; and
- (b) h_p is undersubscribed, or ranks a_i higher than a member of $M(h_p)$.

A couple (a_i, a_j) and an acceptable pair of distinct hospitals h_p and h_q *block* M if

- (c) a_i and a_j are unmatched, or (a_i, a_j) prefers (h_p, h_q) to $(M(a_i), M(a_j))$; and
- (d) h_p is undersubscribed, or $h_p = M(a_i)$, or h_p ranks a_i higher than a member of $M(h_p)$; and
- (e) h_q is undersubscribed, or $h_q = M(a_j)$, or h_q ranks a_j higher than a member of $M(h_q)$.

We say that a couple (a_i, a_j) and a hospital h_p , acceptable to both a_i and a_j , *block* M if

- (f) a_i and a_j are unmatched, or (a_i, a_j) prefers (h_p, h_p) to $(M(a_i), M(a_j))$; and
- (g) either
 - (i) h_p has at least two free places in M ; or
 - (ii) h_p has one free place in M , and $h_p \in \{M(a_i), M(a_j)\}$ or *both* a_i and a_j are higher ranked by h_p than a member of $M(h_p)$; or
 - (iii) h_p is full in M and
 1. $h_p \in \{M(a_i), M(a_j)\}$ and both a_i and a_j are higher ranked by h_p than a member of $M(h_p)$; or
 2. both a_i and a_j are higher ranked by h_p than a member a_k of $M(h_p)$, and a_k is a linked applicant whose partner is also in $M(h_p)$; or

¹In the model described in [4] the applicants are ranked in the same way by all hospital according to a 'master list', and the joint applications of the couples are derived from the individual preferences by the members of the couple in a specific way. Our definition presented here is an extension of that model as we allow hospitals to rank their applicants differently and we have no restriction on how the joint preference lists of the couples are formed.

3. both a_i and a_j are higher ranked by h_p than at least two members of $M(h_p)$.

The rationale behind this stability definition is described in [4] in detail². We do not say that this is the only reasonable stability definition. However, as it is described in [4] and as we will also show in subsection 2.3, this stability definition implies a natural preference ordering over the applications by the hospitals, and the choice functions by the hospitals over the set of applications can be derived from these preferences in a responsive way. This is the key fact that ensures that this problem can be transformed to the ISAC model.

2.2 Integral stable allocation problem with contributions

The stable allocation problem with contributions can be defined for hypergraphs as follows. Suppose that we are given a hypergraph $H(V, E)$ and for each vertex $v \in V(H)$ a strict preference order over the edges incident with v , that corresponds to the preferences of the players over the *contracts* in which they can be involved, where $e <_v f$ denotes that player v prefers contract f to e . Furthermore, we introduce a contribution vector $r_e : V(H) \rightarrow \mathbb{R}_+$ for each edge e of the hypergraph that represents the *contributions* of the agents in contract e . We assume that $v \in e$ if and only if $r_e(v) > 0$, that is when agent v can contribute to contract e . Suppose that we are given nonnegative *bounds* on the vertices $b : V(H) \rightarrow \mathbb{R}_+$ and nonnegative *capacities* on the edges $c : E(H) \rightarrow \mathbb{R}_+$. A nonnegative function x on the edges is an *allocation* if $x(e) \leq c(e)$ for every edge e and $\sum_{e:v \in e} x(e)r_e(v) \leq b(v)$ for every vertex v . An allocation is *stable* if every *unsaturated* edge e (i.e., every edge e with $x(e) < c(e)$) contains a vertex v such that $\sum_{f:v \in f, e <_v f} x(f)r_f(v) = b(v)$. If every bound, capacity and contribution is integral and the problem is to find an integral stable allocation x then we refer to this problem as the *integral stable allocation problem with contributions* (ISAC). Finally, if we allow unit contributions only then we get the *integral stable allocation problem* (ISA).

Biró and Fleiner [3] showed that every stable allocation problem with contributions has a solution that can be obtained by the extended Scarf algorithm. Furthermore, if we have no capacities on the edges then the existence of a stable allocation is guaranteed by the original Scarf lemma, and one stable allocation can be obtained by the original Scarf algorithm.

The exact description of the Scarf algorithm can be found in [18] and its extension for ISAC is described in [3]. Here we only want to highlight some important facts of the Scarf algorithm (and its extension). The algorithm starts with perturbing the capacities (and also the preference matrix if there are ties in the preferences, which is not the case in our particular setting). Based on the perturbation, the Scarf algorithm is deterministic. It takes a well-defined pivot step in each round, that can be implemented efficiently regarding the total number of edges in the hypergraph (i.e. the number of possible contracts). However, we do not have clear ideas on the number of pivot steps the Scarf algorithm should take before termination, we are not aware of results giving lower or upper bounds on that. Some related open questions are listed in [3] and also in the last section of this paper.

2.3 Solving HRC with the Scarf algorithm

First we consider the case when no combined applications are allowed.

Theorem 1 *An instance of HRC with no combined applications can be reduced to an instance of ISA, where the set of stable matchings in the former problem is in one to one correspondence with the set of integral stable allocations in the latter.*

²Here we only note that this particular stability definition reflects the aim of the matching scheme coordinators to have cutoff scores announced for every hospital. A single application is rejected if the resident does not achieve the cutoff score. Similarly, a joint application is rejected if either of the residents in the couple does not achieve the cutoff score at a hospital. A combined application is rejected if the worst resident of the couple either does not achieve the cutoff score or if she/he is better than one admitted single candidate only. Therefore the fairness of the allocation can be verified relatively easily, just like in many higher education matching schemes that use cutoff scores, such as the Hungarian, Irish and Spanish systems, see [6].

PROOF: Suppose that we have an instance of HRC, as described in subsection 2.1. We create an instance of ISA with no edge capacities as follows. Vertex set V of hypergraph $H(V, E)$ represents the agents, namely the single applicants, couples and hospitals. The edges E of $H(V, E)$ correspond to the applications (i.e. possible contracts). For a single application let the contribution of both the resident and the hospital in this contract be one, and similarly, in a joint application let the contributions of the couple and the two hospitals be one each. The capacities of the hospitals are the same as the capacities of the corresponding vertices, the other vertices have unit capacities and no capacity restriction is needed for the edges.

With regard to the vertices representing the single residents and the couples, let their preferences over the edges be the same as the preferences of the single residents and couples over their applications. Regarding the vertices representing the hospitals, the definition of stability described in HRC implies the following preferences by hospitals over the applications. When hospital h_p receives single applications and joint applications only, then it orders these applications according to the residents who are applying to its positions. This may involve indifferences when the same resident is applying to the hospital in different joint applications, but we can suppose that in this case the hospital breaks its ties according to the couple's preferences.

It is now immediate to see the one-to-one correspondence between the set of stable matchings of the HRC instance and the set of integral stable allocations of the ISA instance: an application is accepted if and only if the corresponding edge has unit value in the stable allocation.

□

Theorem 2 *An instance of HRC can be transformed to an instance of ISAC, where the integral stable allocations for the latter problem correspond to the stable matchings for the former one.*

PROOF: Suppose that we have an instance of HRC, as described in subsection 2.1, we create an instance of ISAC as follows. In addition to the proof of Theorem 1, here we need to adjust our construction by accommodating the combined applications. A combined application is also represented by an edge in the hypergraph, where the contribution of the couple is one and the contribution of the hospital is two, since the couple applies for two positions at the hospital.

Regarding the preferences of the hospitals over the applications, the stability definition of HRC implies the following rankings. Suppose first that hospital h_p has to decide between two combined applications, say $[(a_i, a_j) \rightarrow (h_p, h_p)]$ and $[(a_k, a_l) \rightarrow (h_p, h_p)]$, where a_j is higher ranked than a_i and a_l is higher ranked than a_k . In this case h_p will decide according to the worst candidates of these couples, so it would prefer the first application if and only if a_i is higher ranked than a_k . Now, let us consider how the hospital would choose between a resident a_i applying in a single or joint application and a pair of residents (a_j, a_k) applying in a combined application, where a_k is higher ranked by h_p than a_j . Again we suppose that h_p will decide according to its ranking on the weakest candidates, it will prefer the application involving a_i if and only if a_i is higher ranked than a_j by h_p . Again, it is possible that indifferences occur in this preference order between a joint application and a combined application from the same couple, but as previously we suppose that the hospital breaks these ties according to the couple's preferences.

It is again easy to see that an integral stable allocation of the ISAC instance corresponds to a stable matching of the HRC instance, where the edges with unit values represent the accepted applications, since a blocking coalition for the HRC instance would be a blocking edge in the ISAC instance. □

Note that the reverse of the above theorem is not true. A stable matching for a HRC instance might not be an integral stable allocation for the corresponding ISAC instance. This is because in the HRC stability definition it is possible that a combined application is rejected as there is only one place left at a hospital, but a less preferred single or joint application is accepted for that last place. This kind of stable matching cannot be translated to a stable allocation as the rejected combined application would emerge as a blocking edge.

When combined applications are allowed then there may be different reasonable stability concepts. Our transformation described in Theorem 2 works as long as the hospitals can rank the applications, and

from these strict preferences we may derive the hospital choice functions over the set of applications in a quota-responsive way. The latter means that from a set of applications a hospital accepts the best applications one by one, as long as the acceptance of an application does not violate its quota.

Finally we note that if every hospital has one position only then we get a simple stable matching problem for a hypergraph, which is also known as a (*hedonic*) *coalition formation game*. It is worth mentioning that the NP-hardness results [14], [4] hold for this simple case as well. Let us now illustrate the usage of the Scarf algorithm for a particular instance of HRC of this simple kind.

Example

This example was given by Biró, Irving and Schlotter [4] as a difficult instance, for which most heuristics currently used in real applications would fail to find the unique stable solution. Suppose that we have eight residents, comprising three couples (a_1, a_5) , (a_2, a_4) and (a_6, a_8) together with two single applicants a_3 and a_7 . There are eight hospitals, h_1, \dots, h_8 , each with just one post. Suppose that the residents are ordered in the same way by every hospital, according to their indices (a_1 best, a_8 worst), and the individual and joint preference lists of the residents are as follows.

$$\begin{array}{l} a_3 : \quad h_1 \quad h_5 \\ a_7 : \quad h_6 \quad h_8 \\ (a_1, a_5) : (h_1, h_2) \quad (h_3, h_6) \\ (a_2, a_4) : (h_4, h_5) \quad (h_1, h_2) \quad (h_3, h_7) \\ (a_6, a_8) : (h_6, h_8) \end{array}$$

We can describe this problem as a stable matching problem for a hypergraph, where the individual applicants, the couples and the hospitals are represented by vertices, and each hyperedge corresponds to an application (either to an individual application or to a joint application).

To avoid degeneracy and ensure a deterministic invocation of the Scarf algorithm, vector b representing the vertex-bounds should be perturbed. We tried to solve the above problem with the Scarf algorithm by using two different perturbations. First we set $\tilde{b}_i = b_i + \varepsilon_i = b_i + 1/p_{101-i}$, where p_i is the i -th prime number, and we obtained the following half-integral solution:

$$x([a_3 \rightarrow h_3]) = 0, x([a_3 \rightarrow h_5]) = 1, x([a_7 \rightarrow h_6]) = \frac{1}{2}, x([a_7 \rightarrow h_8]) = \frac{1}{2}, x([(a_1, a_5) \rightarrow (h_1, h_2)]) = 1, x([(a_1, a_5) \rightarrow (h_3, h_6)]) = 0, x([(a_2, a_4) \rightarrow (h_4, h_5)]) = 0, x([(a_2, a_4) \rightarrow (h_1, h_2)]) = 0, x([(a_2, a_4) \rightarrow (h_3, h_7)]) = 1, x([(a_6, a_8) \rightarrow (h_6, h_8)]) = \frac{1}{2}.$$

However, by setting $\tilde{b}_i = b_i + \varepsilon_i = b_i + 1/p_{101+i}$, where p_i is the i -th prime number, we obtained

$$x([a_3 \rightarrow h_3]) = 0, x([a_3 \rightarrow h_5]) = 1, x([a_7 \rightarrow h_6]) = 0, x([a_7 \rightarrow h_8]) = 1, x([(a_1, a_5) \rightarrow (h_1, h_2)]) = 0, x([(a_1, a_5) \rightarrow (h_3, h_6)]) = 1, x([(a_2, a_4) \rightarrow (h_4, h_5)]) = 0, x([(a_2, a_4) \rightarrow (h_1, h_2)]) = 1, x([(a_2, a_4) \rightarrow (h_3, h_7)]) = 0, x([(a_6, a_8) \rightarrow (h_6, h_8)]) = 0.$$

which corresponds to the unique stable matching for this instance, namely

$$M = \{\langle a_1, h_3 \rangle, \langle a_2, h_1 \rangle, \langle a_3, h_5 \rangle, \langle a_4, h_2 \rangle, \langle a_5, h_6 \rangle, \langle a_7, h_8 \rangle\}.$$

Therefore this example also illustrates that different perturbations may result in different stable allocations with the possibility that some of them are integral and some are fractional solutions.

3 Experimental comparison of heuristics

In the redesign of the NRMP, Roth and Peranson [16] constructed a heuristic method that incorporates the admission of couples. Motivated by the redesign of the Scottish resident allocation scheme, Biró, Irving

and Schlotter [4] implemented some more complex heuristics and compared them with some variants of the Roth-Peranson method. Finally, we also implemented the Scarf algorithm, and used it as a new heuristic to solve HRC, as described in the previous section. We conducted simulations on the same instances that Biró, Irving and Schlotter [4] studied. But we shall note that the Scarf algorithm was implemented in a different platform and so we did not compare its running time with the other heuristics studied in [4]. For each run of the Scarf algorithm we used one perturbation only. (It is question for future research whether different perturbations can result in significantly different stable allocations for random or structured instances). The table below summarises our findings.

Algorithm	Number of couples										
	12	25	50	75	100	125	150	175	200	225	250
Roth-Peranson	952	897	701	547	395	277	170	83	41	9	3
Best heuristics in B-I-S	976	958	911	870	811	752	682	546	281	71	10
Scarf (int. solution)	895	813	649	532	426	356	316	261	202	174	158
Scarf half-int. solution	999	997	978	958	918	859	816	777	692	695	588
Scarf frac. solution	105	187	351	468	574	644	684	739	798	826	842
Av. # of frac. weights	3.9	4.8	5.7	6.7	7.6	8.8	10.0	10.8	12.8	13.5	15.7
# of frac. weights = 1	41	61	104	127	119	126	106	114	97	85	69
# of frac. weights = 2	2	9	21	30	36	41	43	43	44	48	41
# of frac. weights = 3	14	14	29	38	38	33	35	44	29	36	22
# of frac. weights = 4	7	18	19	25	40	37	39	38	30	32	41
# of frac. weights = 5	11	19	18	25	33	42	34	30	40	28	30

Table 1: Randomly created instances with couples for 500 residents.

In this experiment there were 500 residents and the proportion of couples varied from 5% to 100%. For each proportion of couples there were 1000 random instances generated and we counted how many instances each variant could solve. We remark that we do not know how many of these random instances were actually solvable. Note also that for the set of heuristics by Biró, Irving and Schlotter the results given in the table are not for a single heuristic, but they are always for the heuristic which performed best on the corresponding set of parameter values.

The table shows that the heuristics by Biró, Irving and Schlotter [4] obtained a much better success ratio than the Roth-Peranson heuristics, especially for high proportions of couples. But surprisingly, the Scarf-algorithm was much better than the others when (almost) all the applicants form couples. As we have already noted, this situation occurs in some applications, such as the Hungarian higher education matching scheme (where many applicants apply for a pair of studies).

For completeness, we also included some statistics for those instances that the extended Scarf algorithm could not solve. In particular, we listed how many times the Scarf heuristic returned half-integer solutions (that may be interpreted as half-time contracts), and for how many instances the solution contained only one, two, three, four or five fractional weights. The latter properties of the fractional solutions indicate how far they were from being integral. In particular, having one single fractional weight in a stable allocation must imply that a combined application has half weight, meaning that these two residents would get two half-time jobs in a hospital. In this case, if we would omit the couple and reduce the capacity of that hospital by one then the remaining matching would be stable for the reduced instance.

Further notes

In this paper we showed how the original Scarf algorithm may be used as a successful heuristic for solving the Hospitals Residents problem with Couples, that is a relevant problem in many practical applications. There are still a number of interesting questions that would be worth investigating in extension to our work presented. How does the Scarf algorithm work for the particular case of HRC? Is there some more

efficient way to run this algorithm for HRC than the original matrix method? What are the possible effects of the perturbations used in the Scarf algorithm? If the Scarf algorithm returns a fractional solution, can we approximate an integral solution from that?

We also believe that the stable allocation problem with contributions may accommodate many further relevant problems, and so the Scarf algorithm and its extensions are worth considering as heuristics for solving other problems as well.

References

- [1] B. Aldershof and O.M. Carducci. Stable matchings with couples. *Discrete Applied Mathematics*, 68:203–207, 1996.
- [2] R. Aharoni and T. Fleiner. On a lemma of Scarf. *J. Combin. Theory Ser. B*, 87(1):72–80, 2003. Dedicated to Crispin St. J. A. Nash-Williams.
- [3] P. Biró and T. Fleiner. Fractional solutions for capacitated NTU-games, with applications to stable matchings. *Discussion Paper, IE-RCRES-HAS, MT-DP 2012/34*, 2012.
- [4] P. Biró and R.W. Irving and I. Schlotter. Stable matching with couples – an empirical study. *ACM Journal of Experimental Algorithmics*, 16, Article No.: 1.2, 2011.
- [5] P. Biró and F. Klijn. Matching with Couples: A Multidisciplinary Survey. *To appear in International Game Theory Review*, 2013.
- [6] P. Biró and S. Kiselgof. College admissions with stable score-limits. *Discussion Paper, IE-RCRES-HAS, MT-DP 2013/6*, 2013.
- [7] D. Gale and L. S. Shapley. College Admissions and the Stability of Marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.
- [8] R.W. Irving. "Matching medical students to pairs of hospitals: a new variation on an old theme. in *Proceedings of ESA '98*, LNCS, vol. 1461: 381–392, 1998.
- [9] B. Klaus and F. Klijn. Stable matchings and preferences of couples. *J. Econom. Theory*, 121(1):75–106, 2005.
- [10] B. Klaus, F. Klijn, and J. Massó. Some things couples always wanted to know about stable matchings (but were afraid to ask). *Review of Economic Design*, 11:175–184, 2007.
- [11] F. Kojima, P. A. Pathak and A. E. Roth. Matching with couples: Stability and incentives in large markets. *Working paper*, 2010.
- [12] D. Marx and I. Schlotter. Stable assignment with couples: parameterized complexity and local search. In *Proceedings of IWPEC 2009: The Fourth International Workshop on Parameterized and Exact Computation*, volume 5917 of *Lecture Notes in Computer Science*, pages 300–311. Springer, 2009.
- [13] E. McDermid and D.F. Manlove. Keeping partners together: Algorithmic results for the hospitals / residents problem with couples. *Journal of Combinatorial Optimization*, 19:279–303, 2010.
- [14] E. Ronn. NP-complete stable matching problems. *Journal of Algorithms*, 11:285–304, 1990.
- [15] A. E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 6(4):991–1016, 1984.
- [16] A. E. Roth and E. Peranson. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, 1999.

- [17] A.E. Roth. Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory*, 36:537–569, 2008.
- [18] H. E. Scarf. The core of an N person game. *Econometrica*, 35:50–69, 1967.
- [19] J. Sethuraman, C-P. Teo, and L. Qian. Many-to-one stable matching: geometry and fairness. *Mathematics of Operations Research*, 31:581–596, 2006.
- [20] P. Biró. University admission practices - Hungary. matching-in-practice.eu, accessed on 23 May 2012
- [21] R. Irving. Matching practices for entry-labor markets - Scotland. matching-in-practice.eu, accessed on 23 May 2012
- [22] http://www.nrmp.org/about_nrmp/how.html (National Resident Matching Program website).
- [23] <http://www.nes.scot.nhs.uk/sfas/> (Scottish Foundation Allocation Scheme website).