Available online at www.sciencedirect.com
ScienceDirect
Operations Research Letters 35 (2007) 392–402
Operations Research Letters
www.elsevier.com/locate/orl

Characterizations of the optimal stable allocation mechanism

Mourad Baïou[a, b, *], Michel Balinski[c]

[a] Laboratoire LIMOS, Université Clermont II, Campus des Cézeax, B.P. 125-63173 Aubière Cedex, France
[b] École Polytechnique, Laboratoire d'Économétrie, France
[c] CNRS and École Polytechnique, Laboratoire d'Économétrie, 1 rue Descartes, 75005 Paris, France

Received 13 November 2005; accepted 5 June 2006
Available online 17 August 2006

Abstract

The stable allocation problem is the generalization of (0,1)-matching problems to the allocation of real numbers (hours or quantities) between two separate sets of agents. The same unique-optimal matching (for one set of agents) is characterized by each of three properties: "efficiency", "monotonicity", and "strategy-proofness".

Keywords: Stable marriage; Stable assignment; Ordinal transportation; University admissions; Two-sided market; Many-to-many matching

1. Introduction

Stable matchings are problems of *assignment*: agents from opposite sets having preferences are *matched* [9,13,12,17]. The *stable allocation problem* [4] is also a two-sided market with distinct sets of agents where each agent has strict preferences over the opposite set. But each agent is endowed with real numbers—quantities or hours of work—and instead of matching (or allocating 0's and 1's) the problem is to *allocate* real numbers. For example, one set of agents consists of workmen each with a number of available hours of work, the other of employers each seeking a number of hours of work. "Stability" asks that no pair of opposite agents can increase their hours "together" either due to unused capacity or by giving up hours with less preferred partners.

This natural generalization of two-sided matching presents genuinely new problems: the existence of stable allocations in the general case, when the data is real but not integer-valued, does not follow in straightforward fashion from the corresponding arguments used for marriage or university admissions [6]. Alkan [1] studies stable matchings where the preferences of agents are of a more general "revealed" type: Alkan and Gale [2] extend this approach to stable allocations (but call them stable schedules, the term we used initially).

In general, there may exist many stable allocations, so an obvious question becomes: Which should be

* Corresponding author. Laboratoire LIMOS, Université Clermont II, Campus des Cézeax, B.P. 125-63173 Aubière Cedex, France.

E-mail addresses: baiou@cust.univ-bpclermont.fr (M. Baïou), michel.balinski@shs.polytechnique.fr (M. Balinski).

0167-6377/$ - see front matter
doi:10.1016/j.orl.2006.06.004

chosen? It is the central question of this paper. Of course, exactly the same question presents itself for marriage, admissions and many-to-many problems. In the context of admissions an example of Roth [14] shows that there is no satisfactory answer to what is best for universities when it is assumed that they have "responsive preferences" over the set of possible assignments. It has been shown, however, that in admissions and many-to-many problems a satisfactory answer is forthcoming when universities have preferences that are suggested naturally by the structure of the stable matchings themselves [3,5,6]. This paper generalizes these results, showing that when "generalized max–min preferences" are invoked the employee-optimal allocation mechanism may be characterized uniquely by each of three separate properties: "efficiency", "monotonicity", and "strategy-proofness". Examples show that the characterizations do not hold for weaker definitions of preferences.

2. Stable allocations

There are two distinct finite sets of agents, the *row-agents* I ("employees") and the *column-agents* J ("employers"). Each agent has a strict preference order over the agents of the opposite set. They are collectively called Γ. Each employee $i \in I$ has $s(i)$ units of work to offer, each employer $j \in J$ seeks to obtain $d(j)$ units of work, and $\pi(i, j) \geqslant 0$ is the maximum number of units that $i \in I$ may contract with $j \in J$. Accordingly, a *stable allocation problem* is specified by a quadruple (Γ, s, d, π) where Γ is a set of preferences, $s > 0$ a vector of $|I|$ reals, $d > 0$ a vector of $|J|$ reals, and $\pi \geqslant 0$ an $|I|$ by $|J|$ matrix of reals.

Notation. $i' >_j i$ means that agent $j \in J$ prefers i' to i in I, and similarly, $j' >_i j$ means that agent $i \in I$ prefers j' to j in J. If either $i \in I$ or $j \in J$ refuses to work with the other, then $\pi(i, j) = 0$. The set $(i, j^{>}) \stackrel{\text{def}}{=} \{(i, l) : l >_i j\}$ identifies all agents $l \in J$ that are strictly preferred by row-agent i to column-agent j; and $(i, j^{\geqslant}) \stackrel{\text{def}}{=} \{(i, l) : l \geqslant_i j\}$ all that are strictly preferred as well ascolumn-agent j. The sets $(i^{>}, j)$ and $(i^{\geqslant}, j)$ are defined similarly. In general, if T is a set, $(r, T) \stackrel{\text{def}}{=} \{(r, t) : t \in T\}$, and similarly for

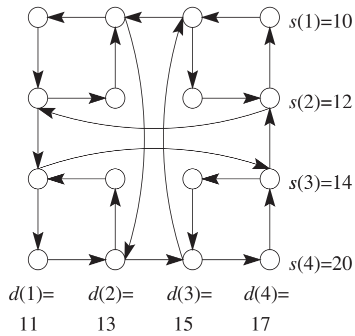

$$y = \begin{pmatrix} 10 & & & \\ 1 & 11 & & \\ & & 14 & \\ & 2 & 1 & 17 \end{pmatrix}$$

$$z = \begin{pmatrix} & 73 & & \\ 11 & & & 1 \\ & & & 14 \\ & 6 & 12 & 2 \end{pmatrix}$$

Fig. 1. An allocation problem (no upper bounds π).

(T, r); moreover, if $y(t)$, $t \in T$, is a real number, then $y(T) \stackrel{\text{def}}{=} \sum_{t \in T} y(t)$.

An *allocation* $x = (x(i, j))$ of a problem (Γ, s, d, π) is a set of real-valued numbers satisfying

$$x(i, J) \leqslant s(i) \quad \text{all } i \in I,$$
$$x(I, j) \leqslant d(j) \quad \text{all } j \in J,$$
$$0 \leqslant x(i, j) \leqslant \pi(i, j) \quad \text{all } (i, j) \in I \times J$$

called, respectively, the *row*, the *column* and the *entry* constraints.

An allocation x is *stable* if for every $(i, j) \in I \times J$,

$$x(i, j) < \pi(i, j) \text{ implies}$$
$$x(i, j^{\geqslant}) = s(i) \text{ or } x(i^{\geqslant}, j) = d(j).$$

If for some (k, l) this condition fails, then (k, l) *blocks* x: agents $k \in I$ and $l \in J$ may together, ignoring others, improve the allocation for themselves. Specifically, the value of $x(k, l)$ may be increased by $\delta > 0$, with $x(k, j) > 0$ for some $j <_k l$ decreased by δ (or $x(k, J) < s(k)$) and $x(i, l) > 0$ for some $i <_l k$ decreased by δ (or $x(I, l) < d(l)$). Otherwise, (k, l) is *stable* for x. In particular, if either $x(k, l) = \pi(k, l)$ or $x(k, l^{\geqslant}) = s(k)$ then (k, l) is *row-stable*; and if either $x(k, l) = \pi(k, l)$ or $x(k^{\geqslant}, l) = d(l)$ then (k, l) is *column-stable*—so a node may be both row- and column-stable.

The "preference graph" of Fig. 1 gives an example with four row-agents and four column-agents, the supply $s(i)$ is attached to each row i, the demand $d(j)$ to each column j, and there are no upper bounds.