ARTICLE IN PRESS

Mathematical Social Sciences ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/econbase

Computational aspects of assigning agents to a line☆

Haris Aziz[a,b], Jens Leth Hougaard[c], Juan D. Moreno-Ternero[d,e], Lars Peter Østerdal[f,*]

[a] *Data61, Australia*
[b] *University of New South Wales, Australia*
[c] *University of Copenhagen, Denmark*
[d] *Universidad Pablo de Olavide, Spain*
[e] *CORE, Université catholique de Louvain, Belgium*
[f] *Copenhagen Business School, Denmark*

HIGHLIGHTS

- We consider the problem of assigning agents to slots on a line.
- We introduce an approach to compute aggregate gap-minimizing assignments.
- We also extend the approach to gap-egalitarian and probabilistic assignments.
- The approach relies on an algorithm which is faster than general purpose algorithms.

ARTICLE INFO

Article history:
Available online xxxx

ABSTRACT

We consider the problem of assigning agents to slots on a line, where only one agent can be served at a slot and each agent prefers to be served as close as possible to his target. We introduce a general approach to compute aggregate gap-minimizing assignments, as well as gap-egalitarian assignments. The approach relies on an algorithm which is shown to be faster than general purpose algorithms for the assignment problem. We also extend the approach to probabilistic assignments and explore the computational features of existing, as well as new, methods for this setting.

1. Introduction

An assignment problem refers to the broad situation in which a set of objects are to be allocated among a group of agents, and each agent is to receive exactly one object. The interest on problems of this sort, which abound in real life, ranges from ancient writings to modern scientific research in different fields (Hylland and Zeckhauser, 1979; Hofstee, 1990; Bogomolnaia and Moulin, 2001; Burkhard et al., 2009).

In a recent paper, Hougaard et al. (2014) introduce and analyze a specific assignment problem, known as the problem of *assigning agents to a line*. In such a problem, each agent has a preferred slot (target) and wants to be served as close as possible to it. Each slot can serve only one agent. The number of slots is at least the number of agents but can be arbitrarily larger than the number of agents.

Hougaard et al. (2014) focus on the notion of (aggregate) *gap minimization* for such a problem, i.e., minimizing the sum of the distances. More precisely, they provide a direct method for testing if a given deterministic assignment is aggregate gap-minimizing, and make use of it to propose an aggregate gap-minimizing modification of the classic *random priority* method to solve this class of problems. It is shown that aggregate gap minimization is incompatible with sd-no-envy and sd-strategy-proofness. Moreover, it is shown that the results extend to more general preference structures. However, the computational aspects of such methods are yet to be explored, and we aim to do so in this paper. Nevertheless, the aims of this paper, which we enumerate next, will go beyond

DOI of original article: http://dx.doi.org/10.1016/j.mathsocsci.2016.12.004.

☆ Paper was prepared for Herve Moulin's Festschrift. We thank him for being a constant source of inspiration for our research. We also thank Maurice Queyranne and two anonymous referees for helpful comments and suggestions on this paper. Financial support from the Danish Council for Independent Research (Grant ID: DFF-1327-0009) and the Spanish Ministry of Economy and Competitiveness (ECO2014-57413-P) is gratefully acknowledged.

* Corresponding author.
E-mail address: lpo.eco@cbs.dk (L.P. Østerdal).

http://dx.doi.org/10.1016/j.mathsocsci.2017.02.004
0165-4896/

studying the computational aspects of the methods developed in Hougaard et al. (2014).

First, we are concerned with the more general problem of identifying (rather than just testing) aggregate gap-minimizing assignments. We introduce an algorithm, dubbed as the *Neat Shifting Algorithm*, to compute aggregate gap-minimizing assignments, which is shown to be faster than general purpose algorithms for the assignment problem. The algorithm relies on the concept of *neatness*, which refers to target-ordered assignments in which all agents with the same target are placed next to each other.

Second, we are also concerned with identifying *gap-egalitarian* assignments, i.e., assignments lexicographically minimizing the distribution of gaps.[1] We show that gap-egalitarian and aggregate gap-minimizing assignments may not coincide; in particular, a gap-egalitarian assignment may not be aggregate gap-minimizing and an aggregate gap-minimizing assignment may not be gap-egalitarian. We also show that, whereas neatness is a necessary condition for gap-egalitarian assignments, there always exists at least one aggregate gap-minimizing assignment that is neat. Furthermore, we show that a suitable adaptation of the *Neat Shifting Algorithm* can be used to yield gap-egalitarian assignments.

Third, when focusing on probabilistic assignments, which arise when randomizing among outcomes, we present new methods, each having their own merits.

The rest of the paper is organized as follows. In Section 2, we discuss some related literature to this paper. In Section 3, we introduce the model and preliminary concepts. In Section 4, we focus on aggregate gap-minimizing assignments and introduce the Neat Shifting Algorithm. In Section 5, we extend the analysis to probabilistic assignments and present aggregate gap-minimizing modifications of the canonical random priority solution. In Section 6, we focus on gap-egalitarian assignments. Section 7 concludes the paper.

2. Related work

As mentioned above, the closest work to this paper is Hougaard et al. (2014) in which the problem of assigning agents to a line is considered. Therein, a direct method for testing if a given deterministic assignment is aggregate gap-minimizing is provided. Chun and Park (forthcoming) consider the same model but assume cardinal (and not just ordinal) preferences. More precisely, they assume that each agent's utility is equal to the amount of monetary transfer minus the distance from the target to his assigned slot. They look at aggregate gap-minimizing as well as gap-egalitarian assignments. They obtain some characterizations but, in contrast to Hougaard et al. (2014), they assume that the number of slots is equal to the number of agents, which greatly simplifies the structure of the assignments as well as the complexity of the problem.[2]

In related work, Procaccia and Tennenholtz (2013) consider a similar problem in the context of 'facility location' and pursue the goals of minimizing the aggregate gap, as well as minimizing the maximum gap. Instead of considering optimal solutions, the focus is on *strategy-proof* mechanisms that provide good approximations of the optimal solutions.

The problem of assigning agents to a line is a slightly different version of the model introduced by Bogomolnaia and Moulin (2001), which could be considered as one of the most influential papers on assignment problems within the economic literature. We not only allow for weak preferences, as Katta and Sethuraman (2006), but we actually assume that preferences are *single-peaked*, as Kasajima (2013) does, and symmetric (to both sides of the peak).

Assignment problems have long been analyzed within the economics literature, mostly focusing on issues of efficiency, incentive compatibility and fairness. Hylland and Zeckhauser (1979) proposed an algorithm, based on market-clearing prices, for allocating individuals to positions with limited capacities. The algorithm guarantees (ex-ante) efficiency, and envy-freeness, but fails to be strategy-proof. As a matter of fact, Zhou (1990) showed that no solution in such a setting satisfies strategy-proofness, (ex-ante) efficiency, and a notion of fairness weaker than envy-freeness. Bogomolnaia and Moulin (2001) restricted attention to strict preferences and ordinal solutions and introduced the notion of *sd-efficiency*, which states that a probabilistic assignment is not stochastically dominated with respect to individual preferences over certain objects. They characterized all sd-efficient assignments and showed that sd-efficiency is incompatible with sd-strategy-proofness, and equal treatment of equals.[3] They also showed that, for more than three agents, no solution satisfies sd-envy-freeness and sd-strategy-proofness together, along with equal treatment of equals.

Previous papers in the economics literature have largely ignored computational aspects. The literature on operations research, computer science and artificial intelligence, on the other hand, has covered various computational aspects of assignment problems (Burkhard et al., 2009; Bouveret et al., 2016; Nguyen et al., 2014; Lang and Rothe, 2015). When the number of items to be allocated to the agents is equal to the number of agents, standard maximum weight matching algorithms can be used to compute allocations that maximize total welfare. For maximizing egalitarian welfare, a perfect matching algorithm can be used to compute a maximum egalitarian assignment. However, the problem of maximizing egalitarian welfare is NP-complete if the number of items is more than the number of agents and agents can get multiple items (Demko and Hill, 1988). Aziz et al. (2016) proved that testing Pareto optimality is coNP-complete under additive preferences even if each agent is to be allocated two items. Aziz et al. (2013) presented a general algorithm to compute and test Pareto optimal outcomes in discrete allocation and coalition formation settings with ordinal preferences. More recently, Damamme et al. (2015) investigated the power of dynamics based on rational bilateral deals (swaps) in various settings including ours. Among other things, they prove NP-hardness of deciding whether a utilitarian or egalitarian allocation is reachable. Now, as they acknowledge, our model of assigning agents to slots on a line induces, by default, a more restrictive domain, as the notion of distance is symmetric.

3. Preliminaries

We consider the model introduced by Hougaard et al. (2014). Imagine a facility with a fixed service capacity that can serve one agent at each (equidistant) *slot*. Agents are labeled by letters, $A, B, \ldots$ with generic elements i and j. The number of slots is infinite and identified with the integer numbers. Each agent i has a preferred slot t_i which we refer to as the agent's *target*. We label the agents so that $t_A \leq t_B \leq \cdots$. A *problem of assigning agents to a line* (in short, a *problem*), consists of a set of agents and the list of their corresponding targets (i.e., slots at which they would like to be served).

1 More precisely, we consider first assignments in which the largest gap is as small as possible. Among those, we consider the assignments in which the second largest gap is as small as possible, etc.

2 The computational aspects of the methods introduced in those papers are not considered therein.

3 Abdulkadiroğlu and Sönmez (2003) provided an alternative characterization of sd-efficiency.

Please cite this article in press as: Aziz, H., et al., Computational aspects of assigning agents to a line. Mathematical Social Sciences (2017), http://dx.doi.org/10.1016/j.mathsocsci.2017.02.004