Cechlárová, K., Eirinakis, P., Fleiner, T., Dimitrios Magos, D., Manlove, D., Mourtos, I., Oceláková, E., & Rastegari, B. (2016). Pareto Optimal Matchings in Many-to-Many Markets with Ties. *Theory of Computing Systems*, 59(4), 700-721.
<https://doi.org/10.1007/s00224-016-9677-1>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1007/s00224-016-9677-1](https://doi.org/10.1007/s00224-016-9677-1)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Springer at <https://doi.org/10.1007/s00224-016-9677-1> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Pareto Optimal Matchings in Many-to-Many Markets with Ties

Katarína Cechlárová¹ · Pavlos Eirinakis² · Tamás Fleiner³ · Dimitrios Magos⁴ · David Manlove⁵ · Ioannis Mourtos² · Eva Oceláková¹ · Baharak Rastegari⁵

Published online: 30 April 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract We consider Pareto optimal matchings (POMs) in a many-to-many market of applicants and courses where applicants have preferences, which may include ties, over individual courses and lexicographic preferences over sets of courses. Since this is the most general setting examined so far in the literature, our work unifies and generalizes several known results. Specifically, we characterize POMs and introduce the *Generalized Serial Dictatorship Mechanism with Ties (GSDT)* that effectively handles ties via properties of network flows. We show that GSDT can generate all POMs using different priority orderings over the applicants, but it satisfies truthfulness only for certain such orderings. This shortcoming is not specific to our mechanism; we show that any mechanism generating all POMs in our setting is prone to strategic manipulation. This is in contrast to the one-to-one case (with or without ties), for which truthful mechanisms generating all POMs do exist.

Keywords Pareto optimality · Many-to-many matching · Serial dictatorship · Truthfulness

✉ Baharak Rastegari
baharak.rastegari@glasgow.ac.uk

¹ Institute of Mathematics, Faculty of Science, P.J. Šafárik University, Košice, Slovakia

² Department of Management Science and Technology, Athens University of Economics and Business, Athens, Greece

³ Budapest University of Technology and Economics, Magyar tudósok körútja and MTA-ELTE Egerváry Research Group, Budapest, Hungary

⁴ Department of Informatics, Technological Educational Institute of Athens, Egaleo, Greece

⁵ School of Computing Science, University of Glasgow, Glasgow, UK

1 Introduction

We study a many-to-many matching market that involves two finite disjoint sets, a set of applicants and a set of courses. Each applicant finds a subset of courses acceptable and has a preference ordering, not necessarily strict, over these courses. Courses do not have preferences. Moreover, each applicant has a quota on the number of courses she can attend, while each course has a quota on the number of applicants it can admit.

A matching is a set of applicant-course pairs such that each applicant is paired only with acceptable courses and the quotas associated with the applicants and the courses are respected. The problem of finding an “optimal” matching given the above market is called the *Course Allocation problem* (CA). Although various optimality criteria exist, *Pareto optimality* (or *Pareto efficiency*) remains the most popular one (see, e.g., [1, 2, 8, 9, 19]). Pareto optimality is a fundamental concept that economists regard as a minimal requirement for a “reasonable” outcome of a mechanism. A matching is a Pareto optimal matching (POM) if there is no other matching in which no applicant is worse off and at least one applicant is better off. Our work examines Pareto optimal many-to-many matchings in the setting where applicants’ preferences may include ties. In the special case where each applicant and course has quota equal to one, our setting reduces to the extensively studied *House Allocation problem* (HA) [1, 15, 21], also known as the *Assignment problem* [4, 12]. Computational aspects of HA have been examined thoroughly [2, 18] and particularly for the case where applicants’ preferences are strict. In [2] the authors provide a characterization of POMs in the case of strict preferences and utilize it in order to construct polynomial-time algorithms for checking whether a given matching is a POM and for finding a POM of maximum size. They also show that any POM in an instance of HA with strict preferences can be obtained through the well-known *Serial Dictatorship Mechanism* (SDM) [1]. SDM is a straightforward greedy algorithm that allocates houses sequentially according to some exogenous priority ordering of the applicants.

Recently, the above results have been extended in two different directions. The first one [17] considers HA in settings where preferences may include ties. Prior to [17], few works in the literature had considered extensions of SDM to such settings. The difficulty regarding ties, observed already in [20], is that the assignments made in the individual steps of the SDM are not unique, and an unsuitable choice may result in an assignment that violates Pareto optimality. To see this, consider a setting with two applicants a_1 and a_2 , and two houses h_1 and h_2 . Assume that a_1 finds both houses acceptable and is indifferent between them, and that a_2 finds only h_1 acceptable. The only Pareto optimal matching for this setting is $\{(a_1, h_2), (a_2, h_1)\}$. Assume that a_1 is served first and that, as both houses are equally acceptable to her, is assigned h_1 (after an arbitrary tie-breaking). Therefore when a_2 ’s turn arrives, there is no house remaining that she finds acceptable, and is hence left unmatched, resulting in a matching that is not Pareto optimal. In [5] and [20] an implicit extension of SDM is provided (in the former case for *dichotomous preferences*, where an applicant’s preference list comprises a single tie containing all acceptable houses), but without an explicit description of an algorithmic procedure. Krysta et al. [17] describe a mechanism called the *Serial Dictatorship Mechanism with Ties* (SDMT) that combines

SDM with the use of augmenting paths to ensure Pareto optimality. SDMT includes an augmentation step, in which applicants already assigned a house may exchange it for another, equally preferred one, to enable another applicant to take a house that is most preferred given the assignments made so far. They also show that any POM in an instance of HA with ties can be obtained by an execution of SDMT and also provide the so-called *Random Serial Dictatorship Mechanism with Ties (RSDMT)* whose (expected) approximation ratio is $\frac{e}{e-1}$ with respect to the maximum-size POM.

The second direction [8] extends the results of [2] to the many-to-many setting (i.e., CA) with strict preferences, while also allowing for a structure of applicant-wise acceptable sets that is more general than the one implied by quotas. Namely, [8] assumes that each applicant selects from a family of course subsets that is downward closed with respect to inclusion. This work provides a characterization of POMs assuming that the preferences of applicants over sets of courses are obtained from their (strict) preferences over individual courses in a lexicographic manner. Using this characterization, it is shown that deciding whether a given matching is a POM can be accomplished in polynomial time. In addition, [8] generalizes SDM to provide the *Generalized Serial Dictatorship Mechanism (GSD)*, which can be used to obtain any POM for CA under strict preferences. The main idea of GSD is to allow each applicant to choose not her most preferred set of courses at once but, instead, only one course at a time (i.e., the most preferred among non-full courses that can be added to the courses already chosen). This result is important as the version of SDM where an applicant chooses immediately her most preferred set of courses cannot obtain all POMs.

Our Contribution In the current work, we combine the directions appearing in [17] and [8] to explore the many-to-many setting in which applicants have preferences, which may include ties, over individual courses. We extend these preferences to sets of courses lexicographically, since lexicographic set preferences naturally describe human behavior [13], they have already been considered in models of exchange of indivisible goods [8, 11] and also possess theoretically interesting properties including responsiveness [16].

We provide a characterization of POMs in this setting, leading to a polynomial-time algorithm for testing whether a given matching is Pareto optimal. We introduce the *Generalized Serial Dictatorship Mechanism with Ties (GSDT)* that generalizes both SDMT and GSD. SDM assumes a priority ordering over the applicants, according to which applicants are served one by one by the mechanism. Since in our setting applicants can be assigned more than one course, each applicant can return to the ordering several times (up to her quota), each time choosing just one course. The idea of using augmenting paths [17] has to be employed carefully to ensure that during course shuffling no applicant replaces a previously assigned course for a less preferred one. To achieve this, we utilize methods and properties of network flows. Although we prove that GSDT can generate all POMs using different priority orderings over applicants, we also observe that some of the priority orderings guarantee truthfulness whereas some others do not. That is, there may exist priority orderings for which some applicant benefits from misrepresenting her preferences. This is in contrast to SDM and SDMT in the one-to-one case in the sense that all executions of

these mechanisms induce truthfulness. This shortcoming however is not specific to our mechanism, since we establish that any mechanism generating all POMs is prone to strategic manipulation by one or more applicants.

For an extended abstract containing the results presented in this paper, see [7].

Organization of the Paper In Section 2 we define our notation and terminology. The characterization is provided in Section 3, while GSDT is presented in Section 4. A discussion on applicants' incentives in GSDT is provided in Section 5. Avenues for future research are discussed in Section 6.

2 Preliminary Definitions of Notation and Terminology

Let $A = \{a_1, a_2, \dots, a_{n_1}\}$ be the set of applicants, $C = \{c_1, c_2, \dots, c_{n_2}\}$ the set of courses and $[i]$ denote the set $\{1, 2, \dots, i\}$. Each applicant a has a quota $b(a)$ that denotes the maximum number of courses a can accommodate into her schedule, and likewise each course c has a quota $q(c)$ that denotes the maximum number of applicants it can admit. Each applicant finds a subset of courses acceptable and has a transitive and complete preference ordering, not necessarily strict, over these courses. We write $c \succ_a c'$ to denote that applicant a (strictly) prefers course c to course c' , and $c \simeq_a c'$ to denote that a is indifferent between c and c' . We write $c \succeq_a c'$ to denote that a either prefers c to c' or is indifferent between them, and say that a weakly prefers c to c' .

Because of indifference, each applicant divides her acceptable courses into *indifference classes* such that she is indifferent between the courses in the same class and has a strict preference over courses in different classes. Let C_t^a denote the t 'th indifference class, or *tie*, of applicant a where $t \in [n_2]$. The preference list of any applicant a is the tuple of sets C_t^a , i.e., $P(a) = (C_1^a, C_2^a, \dots, C_{n_2}^a)$; we assume that $C_t^a = \emptyset$ implies $C_{t'}^a = \emptyset$ for all $t' > t$. Occasionally we consider $P(a)$ to be a set itself and write $c \in P(a)$ instead of $c \in C_t^a$ for some t . We denote by $\mathcal{P}$ the joint preference profile of all applicants, and by $\mathcal{P}(-a)$ the joint profile of all applicants except a . Under these definitions, an instance of CA is denoted by $I = (A, C, \mathcal{P}, b, q)$. Such an instance appears in Table 1.

A (many-to-many) assignment μ is a subset of $A \times C$. For $a \in A$, $\mu(a) = \{c \in C : (a, c) \in \mu\}$ and for $c \in C$, $\mu(c) = \{a \in A : (a, c) \in \mu\}$. An assignment μ is a *matching* if $\mu(a) \subseteq P(a)$ —and thus μ is individually rational, $|\mu(a)| \leq b(a)$ for each $a \in A$ and $|\mu(c)| \leq q(c)$ for each $c \in C$. We say that a is *exposed* if

Table 1 An instance I of CA

Applicant	Quota	Preference list	Course	Quota
a_1	2	$(\{c_1, c_2\}, \{c_3\}, \emptyset)$	c_1	2
a_2	3	$(\{c_2\}, \{c_1, c_3\}, \emptyset)$	c_2	1
a_3	2	$(\{c_3\}, \{c_2\}, \{c_1\})$	c_3	1

$|\mu(a)| < b(a)$, and is *full* otherwise. Analogous definitions of exposed and full hold for courses.

For an applicant a and a set of courses S , we define the *generalized characteristic vector* $\chi_a(S)$ as the vector $(|S \cap C_1^a|, |S \cap C_2^a|, \dots, |S \cap C_{n_2}^a|)$. We assume that for any two sets of courses S and U , a prefers S to U if and only if $\chi_a(S) >_{lex} \chi_a(U)$, i.e., if and only if there is an indifference class C_t^a such that $|S \cap C_t^a| > |U \cap C_t^a|$ and $|S \cap C_{t'}^a| = |U \cap C_{t'}^a|$ for all $t' < t$. If a neither prefers S to U nor U to S , then she is indifferent between S and U . We write $S \succ_a U$ if a prefers S to U , $S \simeq_a U$ if a is indifferent between S and U , and $S \succeq_a U$ if a weakly prefers S to U .

A matching μ is a *Pareto optimal matching* (POM) if there is no other matching in which some applicant is better off and none is worse off. Formally, μ is Pareto optimal if there is no matching μ' such that $\mu'(a) \succeq_a \mu(a)$ for all $a \in A$, and $\mu'(a') \succ_a \mu(a')$ for some $a' \in A$. If such a μ' exists, we say that μ' *Pareto dominates* μ .

A *deterministic mechanism* ϕ maps an instance to a matching, i.e. $\phi : I \mapsto \mu$ where I is a CA instance and μ is a matching in I . A *randomized mechanism* ϕ maps an instance to a distribution over possible matchings. Applicants' preferences are private knowledge and an applicant may prefer not to reveal her preferences truthfully. A deterministic mechanism is *dominant strategy truthful* (or just *truthful*) if all applicants always find it best to declare their true preferences, no matter what other applicants declare. Formally speaking, for every applicant a and every possible declared preference list $P'(a)$, $\phi(P(a), \mathcal{P}(-a)) \succeq_a \phi(P'(a), \mathcal{P}(-a))$, for all $P(a), \mathcal{P}(-a)$. A randomized mechanism ϕ is *universally truthful* if it is a probability distribution over deterministic truthful mechanisms.

3 Characterizing Pareto Optimal Matchings

Manlove [18, Sec. 6.2.2.1] provided a characterization of Pareto optimal matchings in HA with preferences that may include indifference. He defined three different types of *coalitions* with respect to a given matching such that the existence of any means that a subset of applicants can trade among themselves (possibly using some exposed course) ensuring that, at the end, no one is worse off and at least one applicant is better off. He also showed that if no such coalition exists, then the matching is guaranteed to be Pareto optimal. We show that this characterization extends to the many-to-many setting, although the proof is more complex and involved than in the one-to-one setting. We then utilize the characterization in designing a polynomial-time algorithm for testing whether a given matching is Pareto optimal.

In what follows we assume that in each sequence $\mathcal{C}$ no applicant or course appears more than once.

An *alternating path coalition* with respect to μ comprises a sequence $\mathcal{C} = \langle c_{j_0}, a_{i_0}, c_{j_1}, a_{i_1}, \dots, c_{j_{r-1}}, a_{i_{r-1}}, c_{j_r} \rangle$ where $r \geq 1$, $c_{j_k} \in \mu(a_{i_k})$ ($0 \leq k \leq r-1$), $c_{j_k} \notin \mu(a_{i_{k-1}})$ ($1 \leq k \leq r$), a_{i_0} is full, and c_{j_r} is an exposed course. Furthermore, a_{i_0} prefers c_{j_1} to c_{j_0} and, if $r \geq 2$, a_{i_k} weakly prefers $c_{j_{k+1}}$ to c_{j_k} ($1 \leq k \leq r-1$).

An *augmenting path coalition* with respect to μ comprises a sequence $\mathfrak{C} = \langle a_{i_0}, c_{j_1}, a_{i_1}, \dots, c_{j_{r-1}}, a_{i_{r-1}}, c_{j_r} \rangle$ where $r \geq 1$, $c_{j_k} \in \mu(a_{i_k})$ ($1 \leq k \leq r-1$), $c_{j_k} \notin \mu(a_{i_{k-1}})$ ($1 \leq k \leq r$), a_{i_0} is an exposed applicant, and c_{j_r} is an exposed course. Furthermore, a_{i_0} finds c_{j_1} acceptable and, if $r \geq 2$, a_{i_k} weakly prefers $c_{j_{k+1}}$ to c_{j_k} ($1 \leq k \leq r-1$).

A *cyclic coalition* with respect to μ comprises a sequence $\mathfrak{C} = \langle c_{j_0}, a_{i_0}, c_{j_1}, a_{i_1}, \dots, c_{j_{r-1}}, a_{i_{r-1}} \rangle$ where $r \geq 2$, $c_{j_k} \in \mu(a_{i_k})$ ($0 \leq k \leq r-1$), and $c_{j_k} \notin \mu(a_{i_{k-1}})$ ($1 \leq k \leq r$). Furthermore, a_{i_0} prefers c_{j_1} to c_{j_0} and a_{i_k} weakly prefers $c_{j_{k+1}}$ to c_{j_k} ($1 \leq k \leq r-1$). (All subscripts are taken modulo r when reasoning about cyclic coalitions).

We define an *improving coalition* to be an alternating path coalition, an augmenting path coalition or a cyclic coalition.

Given an improving coalition $\mathfrak{C}$, the matching

$$\mu^{\mathfrak{C}} = (\mu \setminus \{(a_{i_k}, c_{j_k}) : \delta \leq k \leq r-1\}) \cup \{(a_{i_k}, c_{j_{k+1}}) : 0 \leq k \leq r-1\} \quad (1)$$

is defined to be the matching obtained from μ by satisfying $\mathfrak{C}$ ($\delta = 1$ in the case that $\mathfrak{C}$ is an augmenting path coalition, otherwise $\delta = 0$).

We will soon show that improving coalitions are at the heart of characterizing Pareto optimal matchings (Theorem 2). The next lemma will come in handy in the proof of Theorems 2 and 10. We say that a sequence of applicants and courses is a *pseudocoalition* if it satisfies all conditions of an improving coalition except that some courses or applicants may appear more than once.

Lemma 1 *Let a matching μ in an instance I of CA admit a pseudocoalition K of length ℓ for some finite ℓ . Then μ admits an improving coalition $\mathfrak{C}$ of length at most ℓ .*

Proof We prove this by induction. Obviously $\ell \geq 2$. For the base case, it is easy to see that K itself is an augmenting path coalition (where $r = 1$) if $\ell = 2$, an alternating path coalition (where $r = 1$) if $\ell = 3$, a cyclic coalition (where $r = 2$) if $\ell = 4$ and K ends with an applicant, and an augmenting path coalition (where $r = 2$) if $\ell = 4$ and K ends with a course. Assume that the claim holds for all pseudocoalitions of length d , $d < \ell$. We show that it also holds for any given pseudocoalition K of length ℓ . In the rest of the proof we show that either K is in fact an improving coalition, or we can derive a shorter pseudocoalition K' from K , hence completing the proof.

The Case for a Repeated Course Assume that a course c appears more than once in K . We consider two different scenarios.

1. If c is not the very first element of K , then sequence K is in the following form where where $c = c_{j_x} = c_{j_y}$.

$$\dots, a_{i_{x-1}}, \overleftarrow{c_{j_x}, a_{i_x}, \dots, c_{j_y}}, \dots$$

In this case, we simply delete the portion of K under the arrow. Note that in the new sequence c_{j_y} appears right after $a_{i_{x-1}}$, but then $c_{j_y} = c_{j_x}$. Hence we have obtained a shorter sequence, and it is easy to verify that this new sequence is a pseudocoalition.

2. If c is the first element of K , then K is in the following form where $c = c_{j_0} = c_{j_y}$.

$$\overleftarrow{c_{j_0}, a_{i_0}, \dots, a_{i_{y-1}}, c_{j_y}, \dots}$$

In this case, we simply only keep the portion of K under the arrow, i.e. our new sequence starts with c_{j_0} and ends with $a_{i_{y-1}}$. The new sequence is shorter and it is easy to verify that it is a pseudocoalition.

The Case for a Repeated Applicant Assume that an applicant a appears more than once in K . We consider the three different possible scenarios.

1. If a is not the very first element of K , nor the last one, then K is in the following form where $a = a_{i_x} = a_{i_y}$.

$$\dots, c_{j_x}, \overleftarrow{a_{i_x}, c_{j_{x+1}}, \dots, a_{i_{y-1}}, c_{i_y}, a_{i_y}, c_{j_{y+1}}, \dots}$$

Note that $c_{j_{x+1}} \succeq_{a_{i_x}} c_{j_x}$ and $c_{j_{y+1}} \succeq_{a_{i_x}} c_{j_y}$ (since $a_{i_x} = a_{j_y}$). We consider two different cases.

- (a) If $c_{j_{x+1}} \succ_{a_{i_x}} c_{j_y}$, then we simply only keep the portion of K under the arrow and add c_{j_y} to the beginning. That is, the new sequence is $K' = \langle c_{j_y}, a_{i_x}, c_{j_{x+1}}, \dots, a_{i_{y-1}} \rangle$. K' is shorter than K and it is easy to verify that that it is a pseudocoalition.
- (b) If $c_{j_y} \succeq_{a_{i_x}} c_{j_{x+1}}$, then $c_{j_{y+1}} \succeq_{a_{i_x}} c_{j_x}$. Then we simply remove the portion of K from $c_{j_{x+1}}$ up to a_{i_y} . That is, the new sequence is $K' = \langle \dots c_{j_x}, a_{i_x}, c_{j_{y+1}} \dots \rangle$ which is shorter than K . Note that if either $c_{j_{x+1}} \succ_{a_{i_x}} c_{j_x}$ or $c_{j_{y+1}} \succ_{a_{i_x}} c_{j_y}$, then $c_{j_{y+1}} \succ_{a_{i_x}} c_{j_x}$. Hence it is easy to verify that K' is a pseudocoalition.
2. If a is the first element of K but not the last one, then K is in the following form where $a = a_{i_0} = a_{i_y}$.

$$a_{i_0}, \overleftarrow{c_{j_1}, \dots, c_{j_y}, a_{i_y}, c_{j_{y+1}}, \dots}$$

Notice that a_0 is exposed and finds $c_{j_{y+1}}$ acceptable and is not matched to it under μ . Hence we simply remove the portion of K under the arrow and get $K' = \langle a_{i_0}, c_{j_{y+1}}, \dots \rangle$. K' is shorter than K and it is easy to verify that it is a pseudocoalition.

3. If a is the last element of K but not the first one, then K is of the following form where $a = a_{i_x} = a_{i_y}$.

$$c_{j_0}, \dots, c_{j_x}, \overleftarrow{a_{i_x}, c_{j_{x+1}}, \dots, a_{i_{y-1}}, c_{j_y}, a_{i_y}}$$

Note that, as $a_{i_x} = a_{i_y}$, $c_{j_{x+1}} \succeq_{a_{i_x}} c_{j_x}$ and $c_{j_0} \succeq_{a_{i_x}} c_{j_y}$. Furthermore, a_{i_x} is not matched to c_{j_0} in μ . We consider two different cases.

- (a) If $c_{j_{x+1}} \succ_{a_{i_x}} c_{j_y}$, then we do as we did in Case 1(a). That is, we only keep the portion of K under the arrow and add c_{j_y} to the beginning; the new sequence is $K' = \langle c_{j_y}, a_{i_x}, c_{j_{x+1}}, \dots, a_{i_{y-1}} \rangle$. K' is shorter than K and it is easy to verify that it is a pseudocoalition.

- (b) If $c_{j_y} \succeq_{a_{i_x}} c_{j_{x+1}}$, then $c_{j_0} \succeq_{a_{i_x}} c_{j_{x+1}} \succeq_{a_{i_x}} c_{j_x}$. Then we simply remove the portion of K from $c_{j_{x+1}}$ until the end of the sequence. That is, the new sequence is $K' = \langle c_{j_0}, \dots, c_{j_x}, a_{i_x} \rangle$. K' is shorter than K and it is easy to verify that it is a pseudocoalition.

Note that we do not need to check the case where K starts and ends with an applicant, since then K would not fit the definition of an improving coalition. $\square$

The following theorem gives a necessary and sufficient condition for a matching to be Pareto optimal.

Theorem 2 *Given a CA instance I , a matching μ is a Pareto optimal matching in I if and only if μ admits no improving coalition.*

Proof Let μ be a Pareto optimal matching in I . Assume to the contrary that μ admits an improving coalition $\mathcal{C}$. It is fairly easy to see that matching $\mu^{\mathcal{C}}$ obtained from μ according to equation (1) Pareto dominates μ , a contradiction.

Conversely, let μ be a matching in I that admits no improving coalition, and suppose to the contrary that it is not Pareto optimal. Therefore, there exists some matching $\mu' \neq \mu$ such that μ' Pareto dominates μ . Let $G_{\mu, \mu'} = \mu \oplus \mu'$ be the graph representing the symmetric difference of μ and μ' . $G_{\mu, \mu'}$ is hence a bipartite graph with applicants in one side and courses in the other; by abusing notation, we may use a or c to refer to a node in $G_{\mu, \mu'}$ corresponding to applicant $a \in A$ or course $c \in C$, respectively. Note that each edge of $G_{\mu, \mu'}$ either belongs to μ , referred to as a μ -edge, or to μ' , referred to as a μ' -edge.

We first provide an intuitive sketch of the rest of the proof. Our goal is to find an improving coalition, hence contradicting the assumption that μ does not admit one. We start with an applicant a_{i_0} who strictly prefers μ' to μ . We choose a μ' -edge (a_{i_0}, c_{j_1}) such that c_{j_1} belongs to the first indifference class where μ' is better than μ for a_{i_0} . Either this edge already represents an improving coalition—which is the case if c_{j_1} is exposed in μ —or there exists a μ -edge incident to c_{j_1} . In this fashion, we continue by adding edges that alternatively belong to μ' (when we are at an applicant node) and μ (when we are at a course node), never passing the same edge twice (by simple book-keeping). Most importantly, we always make sure that when we are at an applicant node, we choose a μ' -edge such that the applicant weakly prefers the course incident to the μ' -edge to the course incident to the precedent μ -edge. We argue that ultimately we either reach an exposed course (which implies that an augmenting or an alternating path coalition exists in μ) or are able to identify a cyclic coalition in μ .

Let us first provide some definitions and facts. Let A' be the set of applicants who prefer μ' to μ ; i.e., $A' = \{a \mid \mu'(a) \succ_a \mu(a)\}$. Moreover, for each $a \in A$, let $\mu'_{\setminus \mu}(a) = \mu'(a) \setminus \mu(a)$ and $\mu_{\setminus \mu'}(a) = \mu(a) \setminus \mu'(a)$. Likewise, for each course $c \in C$, let $\mu_{\setminus \mu'}(c) = \mu(c) \setminus \mu'(c)$ and $\mu'_{\setminus \mu}(c) = \mu'(c) \setminus \mu(c)$. Note that these sets will be altered during the course of the proof. In what follows and in order to simplify presentation, we will say that we remove a μ' -edge (a, c) from $G_{\mu, \mu'}$ to signify that we remove c from $\mu'_{\setminus \mu}(a)$ and remove a from $\mu'_{\setminus \mu}(c)$; similarly, we say that we remove a μ -edge (a, c) from $G_{\mu, \mu'}$ to signify that we remove c from $\mu_{\setminus \mu'}(a)$

and remove a from $\mu_{\setminus\mu'}(c)$. The facts described below follow from the definitions and the assumption that μ' Pareto dominates μ . Let Υ be a set containing applicants (initially, $\Upsilon = \emptyset$); we will explain later what this set will come to contain.

Fact 1: For every course c that is full under μ , it is the case that $|\mu_{\setminus\mu'}(c)| \geq |\mu'_{\setminus\mu}(c)|$.

Fact 2: For every applicant $a \in A' \setminus \Upsilon$, there exists $\ell_a^* \leq n_2$ such that $|C_\ell^a \cap \mu'_{\setminus\mu}(a)| = |C_\ell^a \cap \mu_{\setminus\mu'}(a)|$ for all $\ell < \ell_a^*$, and $|C_{\ell_a^*}^a \cap \mu'_{\setminus\mu}(a)| > |C_{\ell_a^*}^a \cap \mu_{\setminus\mu'}(a)|$. Hence, $C_{\ell_a^*}^a$ denotes the first indifference class of a in which μ' is better than μ for a .

Fact 3: For every applicant $a \in A'$ that is full in μ , there exists ℓ^+ , $\ell_a^* < \ell^+ \leq n_2$, such that $|C_{\ell^+}^a \cap \mu_{\setminus\mu'}(a)| > |C_{\ell^+}^a \cap \mu'_{\setminus\mu}(a)|$.

In what follows, we iteratively remove edges from $G_{\mu,\mu'}$ (so as not to visit an edge twice and maintain the aforementioned facts) and create a corresponding path $\mathfrak{C}$, which we show that in all possible cases implies an improving coalition in μ , thus contradicting our assumption.

As μ' Pareto dominates μ , A' is nonempty. Let a_{i_0} be an applicant in A' . By Fact 2, there exists a course $c_{j_1} \in C_{\ell_{a_{i_0}}^*}^a \cap \mu'_{\setminus\mu}(a_{i_0})$. Hence, we remove edge (a_{i_0}, c_{j_1}) from $G_{\mu,\mu'}$ and add it to $\mathfrak{C}$, i.e., $\mathfrak{C} = \langle a_{i_0}, c_{j_1} \rangle$. If c_{j_1} is exposed, then an augmenting path coalition is implied. More specifically, if a_{i_0} is also exposed in μ , then $\mathfrak{C} = \langle a_{i_0}, c_{j_1} \rangle$ is an augmenting path coalition. Otherwise, if a_{i_0} is full in μ , then it follows from Fact 3 that there exists a course $c_{j_0} \in \mu_{\setminus\mu'}(a_{i_0})$ such that $c_{j_1} \succ_{a_{i_0}} c_{j_0}$. Then $\mathfrak{C} = \langle c_{i_0}, a_{i_0}, c_{j_1} \rangle$ is an alternating path coalition.

If, on the other hand, c_{j_1} is full, we continue our search for an improving coalition in an iterative manner (as follows) until we either reach an exposed course (in which case we show that an augmenting or alternating coalition is found) or revisit an applicant that belongs to A' (in which case we show that a cyclic coalition is found). Note that $a_{i_0} \in \mu'_{\setminus\mu}(c_{j_1})$, hence it follows from Fact 1 that $|\mu_{\setminus\mu'}(c_{j_1})| \geq 1$. At the start of each iteration $k \geq 1$, we have:

$$\mathfrak{C} = \langle a_{i_0}, c_{j_1}, \dots, a_{i_{k-1}}, c_{j_k} \rangle$$

such that by the construction of $\mathfrak{C}$ (through the iterative procedure) every applicant a_{i_x} on $\mathfrak{C}$ weakly prefers the course that follows her on $\mathfrak{C}$ —i.e., $c_{j_{x+1}}$ to which she is matched to in μ' but not in μ —to the course that precedes her—i.e., c_{j_x} to which she is matched to in μ but not in μ' . Moreover, a_{i_0} is either exposed and can accommodate one more course, namely c_{j_1} , or is full and hence there exists a course $c_{j_0} \in \mu_{\setminus\mu'}(a_{i_0})$ such that a_{i_0} strictly prefers c_{j_1} to c_{j_0} . Notice that $\mathfrak{C}$ would imply an improving coalition (using Lemma 1) if c_{j_k} is exposed.

During the iterative procedure, set Υ includes those applicants who strictly prefer the course they are given in $\mathfrak{C}$ under μ' than under μ . That is, $\Upsilon = \{a_{i_k} \in \mathfrak{C} : c_{j_{k+1}} \succ_{a_{i_k}} c_{j_k}\} \cup \{a_{i_0}\}$. Note that applicant a_{i_0} is always included in Υ , since either she is exposed, or $\exists c_{j_0} \in \mu_{\setminus\mu'}(a_{i_0})$ such that $c_{j_1} \succ_{a_{i_0}} c_{j_0}$.

The Iterative Procedure We repeat the following procedure until either the last course on $\mathfrak{C}$, c_{j_k} , is exposed or Case 3 is reached. If c_{j_k} is full, then by Fact 1 there

must exist an applicant $a_{i_k} \in \mu_{\setminus \mu'}(c_{i_k})$; note that $a_{i_k} \neq a_{i_{k-1}}$. Remove (a_{i_k}, c_{j_k}) from $G_{\mu, \mu'}$. Let $C_\ell^{a_{i_k}}$ be the indifference class of a_{i_k} to which c_{j_k} belongs. We consider three different cases.

- **Case 1:** $|C_\ell^{a_{i_k}} \cap \mu'_{\setminus \mu}(a_{i_k})| \geq |C_\ell^{a_{i_k}} \cap \mu_{\setminus \mu'}(a_{i_k})|$. Then there must exist a course $c_{j_{k+1}} \in C_\ell^{a_{i_k}} \cap \mu'_{\setminus \mu}(a_{i_k})$; note that $c_{j_k} \neq c_{j_{k+1}}$. Hence, we remove (a_{i_k}, c_{j_k}) from $G_{\mu, \mu'}$ and append $\langle a_{i_k}, c_{j_{k+1}} \rangle$ to $\mathfrak{C}$.
- **Case 2:** $|C_\ell^{a_{i_k}} \cap \mu'_{\setminus \mu}(a_{i_k})| < |C_\ell^{a_{i_k}} \cap \mu_{\setminus \mu'}(a_{i_k})|$ and $a_{i_k} \notin \Upsilon$. Recall that μ' Pareto dominates μ , hence it must be that $a_{i_k} \in A'$. Therefore it follows from Fact 2 that $\ell_{a_{i_k}}^* < \ell$ and there exists a course $c_{j_{k+1}} \in C_{\ell_{a_{i_k}}^*}^{a_{i_k}} \cap \mu'_{\setminus \mu}(a_{i_k})$; note that $c_{j_k} \neq c_{j_{k+1}}$. Hence, we remove (a_{i_k}, c_{j_k}) from $G_{\mu, \mu'}$ and append $\langle a_{i_k}, c_{j_{k+1}} \rangle$ to $\mathfrak{C}$. Moreover, we add a_{i_k} to Υ .
- **Case 3:** $|C_\ell^{a_{i_k}} \cap \mu'_{\setminus \mu}(a_{i_k})| < |C_\ell^{a_{i_k}} \cap \mu_{\setminus \mu'}(a_{i_k})|$ and $a_{i_k} \in \Upsilon$; note that this case cannot happen when $k = 1$. Let a_{i_z} be the first occurrence of a_{i_k} which has resulted in a_{i_k} being added to Υ ; i.e., $0 \leq z < k$ and $a_{i_z} = a_{i_k}$. Therefore, $\mathfrak{C}$ is of the form $\langle a_{i_0}, c_{j_1}, a_{i_1}, \dots, a_{i_z}, c_{j_{z+1}}, \dots, a_{i_{k-1}}, c_{j_k} \rangle$. Let us consider $\mathfrak{C} = \langle c_{j_k}, a_{i_z}, c_{j_{z+1}}, \dots, a_{i_{k-1}} \rangle$. Since $a_{i_z} = a_{i_k} \in \Upsilon$, $a_{i_z} \in A'$ and, by the construction of $\mathfrak{C}$, $c_{j_{z+1}} \in C_{\ell_{a_{i_k}}^*}^{a_{i_k}}$, it follows from Fact 2 that $\ell_{a_{i_z}}^* < \ell$, which implies (using the assumption of Case 3) that $c_{j_{z+1}} \succ_{i_z} c_{j_k}$.

If at any iteration k we find a c_{j_k} that is exposed in μ or we arrive at Case 3, then an improving coalition is implied (using Lemma 1), which contradicts our assumption on μ . Otherwise (cases 1 and 2), we continue with a new iteration. However, since the number of edges is bounded, this iterative procedure is bound to terminate, either by reaching an exposed course or by reaching Case 3. $\square$

It follows from Theorem 2 that in order to test whether a given matching μ is Pareto optimal, it is sufficient to check whether it precludes any improving coalition. To check this, we construct the *extended envy graph* associated with μ .

Definition 1 The *extended envy graph* $G(\mu) = (V_{G(\mu)}, E_{G(\mu)})$ associated with a matching μ is a weighted digraph with $V_{G(\mu)} = V_A \cup V_C \cup V_\mu$, where $V_C = \{c : c \in C\}$, $V_A = \{a : a \in A\}$ and $V_\mu = \{ac : (a, c) \in \mu\}$, and $E_{G(\mu)} = E_{G(\mu)}^1 \cup E_{G(\mu)}^2 \cup E_{G(\mu)}^3$ where:

- $E_{G(\mu)}^1 = \{(c, a) : a \in A, c \in C, \text{ and } |\mu(c)| < q(c)\} \cup \{(c, ac') : c \in C \text{ and } |\mu(c)| < q(c)\}$, with weight equal to 0;
- $E_{G(\mu)}^2 = \{(a, c) : a \in A, c \in P(a) \setminus \mu(a) \text{ and } |\mu(a)| < b(a)\} \cup \{(a, a'c) : a \in A \setminus \{a'\}, c \in P(a) \setminus \mu(a) \text{ and } |\mu(a)| < b(a)\}$, with weight equal to -1 ; and
- $E_{G(\mu)}^3 = \{(ac, c') : c' \in P(a) \setminus \mu(a) \text{ and } c' \succeq_a c\} \cup \{(ac, a'c') : c' \in P(a) \setminus \mu(a) \text{ and } c' \succeq_a c\}$, with weight equal to 0 if $c' \simeq_a c$ or -1 if $c' \succ_a c$.

That is, $G(\mu)$ includes one node per course, one node per applicant and one node per applicant-course pair that participates in μ . With respect to $E_{G(\mu)}$, a ‘0’ (‘ -1 ’)

arc is an arc with weight 0 (-1). There is a ‘0’ arc leaving from each exposed course node to any other (non-course) node in $G(\mu)$ (arcs in $E_{G(\mu)}^1$). Moreover, there is a -1 arc leaving each exposed applicant node a towards any course node or applicant-course node whose corresponding course is in $P(a) \setminus \mu(a)$ (arcs in $E_{G(\mu)}^2$). Finally, there is an arc leaving each applicant-course node ac towards any course node or applicant-course node whose corresponding course c' is not in $\mu(a)$ and $c' \succeq_a c$. The weight of such an arc is 0 if $c' \simeq_a c$ or -1 if $c' \succ_a c$ (arcs in $E_{G(\mu)}^3$).

The following theorem establishes the connection between Pareto optimality of a given matching and the absence of negative cost cycles in the extended envy graph corresponding to that matching.

Theorem 3 *A matching μ is Pareto optimal if and only if its extended envy graph $G(\mu)$ has no negative cost cycles.*

Proof First we show that if $G(\mu)$ has no negative cost cycle, then μ is Pareto optimal. Assume for a contradiction that μ is not Pareto optimal; this means that there exists an improving coalition $\mathfrak{C}$. We examine each of the three different types of improving coalition separately. Assume that $\mathfrak{C}$ is an alternating path coalition, i.e., $\mathfrak{C} = \langle c_{j_0}, a_{i_0}, c_{j_1}, a_{i_1}, \dots, c_{j_{r-1}}, a_{i_{r-1}}, c_{j_r} \rangle$ where $r \geq 1$, $(a_{i_k}, c_{j_k}) \in \mu$ ($0 \leq k \leq r-1$), $(a_{i_{k-1}}, c_{j_k}) \notin \mu$ ($1 \leq k \leq r$), a_{i_0} is full, and c_{j_r} is an exposed course. Furthermore, $c_{j_1} \succ_{a_{i_0}} c_{j_0}$ and, if $r \geq 2$, $c_{j_{k+1}} \succeq_{a_{i_k}} c_{j_k}$ ($1 \leq k \leq r-1$). By Definition 1, there exists in $G(\mu)$ an arc $(a_{i_k}c_{j_k}, a_{i_{k+1}}c_{j_{k+1}})$ for each $0 \leq k \leq r-2$, an arc $(a_{i_{r-1}}c_{j_{r-1}}, c_{j_r})$, and an arc $(c_{j_r}, a_{i_0}c_{j_0})$, thus creating a cycle K in $G(\mu)$. Since all arcs of K have weight 0 or -1 and $(a_{i_0}c_{j_0}, a_{i_1}c_{j_1})$ has weight -1 (because $c_{j_1} \succ_{a_{i_0}} c_{j_0}$), K is a negative cost cycle, a contradiction. A similar argument can be used for the case where $\mathfrak{C}$ is an augmenting path coalition; our reasoning on the implied cycle would start with a ‘ -1 ’ arc $(a_{i_0}, a_{i_1}c_{j_1})$ and would conclude with a ‘0’ arc (c_{j_r}, a_{i_0}) . Finally, for a cyclic coalition $\mathfrak{C}$, the implied negative cost cycle includes all arcs $(a_{i_k}c_{j_k}, a_{i_{k+1}}c_{j_{k+1}})$ for $0 \leq k \leq r-1$ taken modulo r , where at least one arc has weight -1 (otherwise, $\mathfrak{C}$ would not be a cyclic coalition). In all cases, the existence of an improving coalition implies a negative cost cycle in $G(\mu)$, a contradiction.

Next, we show that if μ is Pareto optimal, $G(\mu)$ has no negative cost cycles. We prove this via contradiction. Let μ be a POM, and assume that $G(\mu)$ has a negative cost cycle. Given that there exists a negative cost cycle in $G(\mu)$, there must also exist one that is simple and of minimal length. Hence, without loss of generality, let K be a minimal (length) simple negative cost cycle in $G(\mu)$. We commence by considering that K contains only nodes in V_μ . We create a sequence $\mathcal{K}$ corresponding to K as follows. Starting from an applicant-course node of K with an outgoing ‘ -1 ’ arc, say $a_{i_0}c_{j_0}$, we first place c_{j_0} in $\mathcal{K}$ followed by a_{i_0} , and continue by repeating this process for all nodes in K . It is then easy to verify, by Definition 1 and the definition of a cyclic coalition, that K corresponds to a sequence $\mathcal{K} = \langle c_{j_0}, a_{i_0}, \dots, c_{j_{r-1}}, a_{i_{r-1}} \rangle$ that satisfies all conditions of a cyclic coalition with the exception that some courses or applicants may appear more than once. But this means that, by Lemma 1, μ admits an improving coalition, which contradicts that μ is a POM.

To conclude the proof, we consider the case where K includes at least one node from V_A or V_C . Note that if K contains a node from V_A , then it also contains a node from V_C . This is because, by Definition 1, the only incoming arcs to an applicant node are of type $E_{G(\mu)}^1$. Therefore K includes at least one course node. We now claim that since K is of minimal length, it will contain only one course node. Assume for a contradiction that K contains more than one course nodes. Since K is a negative cost cycle, it contains at least one ‘ -1 ’ arc which leaves an applicant node a (or an applicant-course node ac'). Let c be the first course node in K that appears after this ‘ -1 ’ arc. Recall that for c to be in K , c is exposed in μ and there exists a ‘ 0 ’ arc from c to any applicant or applicant-course node of $G(\mu)$. Hence, there also exists such an arc from c to a (or to ac'), thus forming a negative cost cycle smaller than K , a contradiction. Note that this also implies that K contains at most one node in V_A . This is because, by Definition 1, the only incoming arcs to nodes in V_A are from nodes in V_C , therefore the claim follows from the facts that c is the only course node in K , and K is a simple cycle. We create a sequence $\mathcal{K}$ corresponding to K as follows. If K contains an applicant node a , we start $\mathcal{K}$ with a . If K contains a course node c , we end $\mathcal{K}$ with c . Any applicant-course node $a'c'$ is handled by placing c' first and then a' . It is now easy to verify, by Definition 1 and the definition of a cyclic coalition, that K corresponds to a sequence $\mathcal{K} = \langle a_{i_0}, c_{j_1}, a_{i_1}, \dots, c_{j_{r-1}}, a_{i_{r-1}}, c_{j_r} \rangle$ (if K contains an applicant node) or a sequence $\mathcal{K} = \langle c_{j_0}, a_{i_0}, \dots, c_{j_{r-1}}, a_{i_{r-1}}, c_{j_r} \rangle$ (if K contains no applicant node), such that $\mathcal{K}$ satisfies all conditions of either an augmenting path coalition (if K contains an applicant node) or an alternating path coalition (if K contains no applicant node), with the exception that some courses or applicants may appear more than once. But this means that, by Lemma 1, μ admits an augmenting path coalition (or an alternating path coalition), which contradicts that μ is a POM. $\square$

Theorem 3 implies that to test whether a matching μ is Pareto optimal or not, it suffices to create the corresponding extended envy graph and check whether it contains a negative cost cycle. It is easy to see that we can create $G(\mu)$ in polynomial time. To test whether $G(\mu)$ admits a negative cycle, we can make use of various algorithms that exist in the literature (see [10] for a survey), e.g. Moore-Bellman-Ford which terminates in $O(|V_{G(\mu)}||E_{G(\mu)}|)$ time. However, as all arcs in $G(\mu)$ are of costs either 0 or -1 and hence integer, we can use a faster algorithm of [14] which, for our setting, terminates in $O(\sqrt{|V_{G(\mu)}|}|E_{G(\mu)}|)$ time.

Corollary 4 *Given a CA instance I and a matching μ , we can check in polynomial time whether μ is Pareto optimal in I .*

4 Constructing Pareto Optimal Matchings

We propose an algorithm for finding a POM in an instance of CA, which is in a certain sense a generalization of Serial Dictatorship thus named *Generalized Serial Dictatorship Mechanism with ties* (GSDT). The algorithm starts by setting the quotas of all applicants to 0 and those of courses at the original values given by q . At each stage i ,

the algorithm selects a single applicant whose original capacity has not been reached, and increases only her capacity by 1. The algorithm terminates after $B = \sum_{a \in A} b(a)$ stages, i.e., once the original capacities of all applicants have been reached. In that respect, the algorithm assumes a ‘multisequence’ $\Sigma = (a^1, a^2, \dots, a^B)$ of applicants such that each applicant a appears $b(a)$ times in Σ ; e.g., for the instance of Table 1 and the sequence $\Sigma = (a_1, a_1, a_2, a_2, a_3, a_2, a_3)$, the vector of capacities evolves as follows:

$$(0, 0, 0), (1, 0, 0), (2, 0, 0), (2, 1, 0), (2, 2, 0), (2, 2, 1), (2, 3, 1), (2, 3, 2).$$

Let us denote the vector of applicants’ capacities in stage i by b^i , i.e., b^0 is the all-zeroes vector and $b^B = b$. Clearly, each stage corresponds to an instance I^i similar to the original instance except for the capacities vector b^i . At each stage i , our algorithm obtains a matching μ^i for the instance I^i . Since the single matching of stage 0, i.e., the empty matching, is a POM in I^0 , the core idea is to modify μ^{i-1} in such way that if μ^{i-1} is a POM with respect to I^{i-1} then μ^i is a POM with respect to I^i . To achieve this, the algorithm relies on the following flow network.

Consider the digraph $D = (V, E)$. Its node set is $V = A \cup T \cup C \cup \{\sigma, \tau\}$ where σ and τ are the source and the sink and vertices in T correspond to the ties in the preference lists of all applicants; i.e., T has a node (a, t) per applicant a and tie t such that $C_t^a \neq \emptyset$. Its arc set is $E = E_1 \cup E_2 \cup E_3 \cup E_4$ where $E_1 = \{(\sigma, a) : a \in A\}$, $E_2 = \{(a, (a, t)) : a \in A, C_t^a \neq \emptyset\}$, $E_3 = \{((a, t), c) : c \in C_t^a\}$ and $E_4 = \{(c, \tau) : c \in C\}$. The graph D for the instance of Table 1 appears in Fig. 1, where an oval encircles all the vertices of T that correspond to the same applicant, i.e., one vertex per tie.

Using digraph $D = (V, E)$, we obtain a flow network N^i at each stage i of the algorithm, i.e., a network corresponding to instance I^i , by appropriately varying the capacities of the arcs. (For an introduction on network flow algorithms see, e.g., [3].) The capacity of each arc in E_3 is always 1 (since each course may be received at most once by each applicant) and the capacity of an arc $e = (c, \tau) \in E_4$ is always $q(c)$. The capacities of all arcs in $E_1 \cup E_2$ are initially 0 and, at stage i , the capacities of only certain arcs associated with applicant a^i are increased by 1. For this reason, for each applicant a we use the variable $curr(a)$ that indicates her ‘active’ tie; initially, $curr(a)$ is set to 1 for all $a \in A$.

In stage i , the algorithm computes a maximum flow f^i whose saturated arcs in E_3 indicate the corresponding matching μ^i . The algorithm starts with $f^0 = 0$ and $\mu^0 = \emptyset$. Let the applicant $a^i \in A$ be a copy of applicant a considered in stage i . The algorithm increases by 1 the capacity of arc $(\sigma, a) \in E_1$ (i.e., the applicant is allowed to receive an additional course). It then examines the tie $curr(a)$ to check whether the additional course can be received from tie $curr(a)$. To do this, the capacity of arc $(a, (a, curr(a))) \in E_2$ is increased by 1. The network in stage i where tie $curr(a^i)$ is examined is denoted by $N^{i, curr(a^i)}$. If there is an augmenting $\sigma - \tau$ path in this network, the algorithm augments the current flow f^{i-1} to obtain f^i , accordingly augments μ^{i-1} to obtain μ^i (i.e., it sets μ^i to the symmetric difference of μ^{i-1} and all pairs (a, c) for which there is an arc $((a, t), c)$ in the augmenting path) and proceeds to the next stage. Otherwise, it decreases the capacity of $(a, (a, curr(a)))$ by 1 (but not the capacity of arc (σ, a)) and it increases $curr(a)$ by 1 to examine the next tie

of a ; if all (non-empty) ties have been examined, the algorithm proceeds to the next stage without augmenting the flow. Note that an augmenting $\sigma - \tau$ path in the network $N^{i, curr(a^i)}$ corresponds to an augmenting path coalition in μ^{i-1} with respect to I^i .

A formal description of GSDT is provided by Algorithm 1, where $w(e)$ denotes the capacity of an arc $e \in E$ and $\oplus$ denotes symmetric difference. Observe that all arcs in E_2 are saturated, except for the arc corresponding to the current applicant and tie, thus any augmenting path has one arc from each of E_1 , E_2 and E_4 and all other arcs from E_3 ; as a consequence, the number of courses each applicant receives at stage i in any tie cannot decrease at any subsequent step. Also, μ^i dominates μ^{i-1} with respect to instance I^i if and only if there is a flow in N^i that saturates all arcs in E_2 .

Algorithm 1 Producing a POM for any instance of CA

Input: an instance I of CA and a multisequence Σ
 $f^0 := 0$; $\mu^0 := \emptyset$;
for each $a \in A$, $curr(a) := 1$;
for $i = 1, 2, \dots, B$ **do**
 {
 consider the applicant $a = a^i$;
 $w(\sigma, a)++$;
 $\mathfrak{P} := \emptyset$;
 while $\mathfrak{P} = \emptyset$ **and** $curr(a) \leq n_2$ **and** $C_{curr(a)}^a \neq \emptyset$ **do**
 {
 $w(a, (a, curr(a)))++$;
 $\mathfrak{P} :=$ augmenting path in $N^{i, curr(a)}$ with respect to f^{i-1} ;
 if $\mathfrak{P} = \emptyset$ **then** { $w(a, (a, curr(a)))--$; $curr(a)++$; }
 }
 if $\mathfrak{P} \neq \emptyset$ **then** { $f^i := f^{i-1} \oplus \mathfrak{P}$;
 $\mu^i := \mu^{i-1} \oplus \{(a, c) : ((a, t), c) \in \mathfrak{P} \text{ for some tie } t\}$; }
 otherwise { $f^i := f^{i-1}$; $\mu^i := \mu^{i-1}$; }
 }
return μ^B ;

To prove the correctness of GSDT, we need two intermediate lemmas. Let $e_t \in \mathbb{R}^{n_2}$ be the vector having 1 at entry t and 0 elsewhere.

Lemma 5 *Let $N^{i,t}$ be the network at stage i while tie t of applicant a^i is examined. Then, there is an augmenting path with respect to f^{i-1} in $N^{i,t}$ if and only if there is a matching μ such that*

$$\chi_a(\mu(a)) = \chi_a(\mu^{i-1}(a)) \text{ for each } a \neq a^i \text{ and } \chi_{a^i}(\mu(a^i)) = \chi_{a^i}(\mu^{i-1}(a^i)) + e_t.$$

Proof Note that the flows f^{i-1} and f^i , corresponding to μ^{i-1} and μ^i respectively, are feasible in N^i . Moreover, f^{i-1} is feasible in $N^{i,t}$ and, if there is an augmenting path with respect to f^{i-1} in $N^{i,t}$, then f^i is feasible in $N^{i,t}$ too.

If there is a path augmenting f^{i-1} by 1 in $N^{i,t}$ thus obtaining f^i , the number of courses assigned per tie and applicant remain identical except for a^i that receives an

Fig. 1 Digraph D for the instance I from Table 1

extra course from tie t . Thus, $\chi_a(\mu(a))$ is identical to $\chi_a(\mu^{i-1}(a))$ for all $a \in A$ except for a^i whose characteristic vector has its t 'th entry increased by 1.

Conversely, if the above equation holds, flow f^i is larger than f^{i-1} in the network $N^{i,t}$. Thus standard network flow theory implies the existence of a path augmenting f^{i-1} . As a side remark, since the network $N^{i,t}$ has similar arc capacities with the network N^{i-1} except for the arcs (σ, a^i) and $(a^i, (a^i, t))$, this augmenting path includes these two arcs. $\square$

Lemma 6 Let $S \succeq_a U$ and $|S| \geq |U|$. If c_S and c_U denote a least preferred course of applicant a in S and U , respectively, then $S \setminus \{c_S\} \succeq_a U \setminus \{c_U\}$.

Proof For convenience, we denote $S \setminus \{c_S\}$ by S' and $U \setminus \{c_U\}$ by U' . Let

$$\chi_a(S) = (s_1, s_2, \dots, s_{n_2}), \quad \chi_a(U) = (u_1, u_2, \dots, u_{n_2})$$

and

$$\chi_a(S') = (s'_1, s'_2, \dots, s'_{n_2}), \quad \chi_a(U') = (u'_1, u'_2, \dots, u'_{n_2}).$$

If $S \simeq_a U$, then $c_S \simeq_a c_U$ and $\chi_a(S) = \chi_a(U)$, therefore $\chi_a(S') = \chi_a(U')$ and $S' \simeq_a U'$. Otherwise, $S \succ_a U$ implies that there is $k \in [n_2]$ such that $s_j = u_j$ for each $j < k$ and $s_k > u_k$.

If there is $j \geq k+1$ with $s_j > 0$, then $c_S \in C_r^a$ with $r \geq k+1$. It follows that $s'_k = s_k > u_k \geq u'_k$ while $s'_j \geq u'_j$ for all $j < k$, thus $\chi_a(S') \succ_{lex} \chi_a(U')$, i.e., $S' \succ_a U'$. The same follows if $s_k \geq u_k + 2$ because then $s'_k \geq s_k - 1 > u_k \geq u'_k$.

It remains to examine the case where $s_j = 0$ for all $j \geq k+1$ and $s_k = u_k + 1$. In this case, $c_S \in C_k^a$ thus $s'_k = s_k - 1 = u_k$, while $|S| \geq |U|$ implies that either $u_j = 0$ for all $j \geq k+1$ or there is a single $k' > k$ such that $u_{k'} = 1$. In the former case $c_U \in C_r^a$ for $r \leq k$ thus $\chi_a(S') \succ_{lex} \chi_a(U')$, whereas in the latter one $u'_k = u_k = s_k - 1 = s'_k$ and $s'_j = u'_j = 0$ for all $j > k$ hence $\chi_a(S') = \chi_a(U')$. $\square$

Theorem 7 For each i , the matching μ^i obtained by GSDT is a POM for instance I^i .

Proof We apply induction on i . Clearly, $\mu^0 = \emptyset$ is the single matching in I^0 and hence a POM in I^0 . We assume that μ^{i-1} is a POM in I^{i-1} and prove that μ^i is a POM in I^i .

Assume to the contrary that μ^i is not a POM in I^i . This implies that there is a matching ξ in I^i that dominates μ^i . Then, for all $a \in A$, $\xi(a) \succeq_a \mu^i(a) \succeq_a \mu^{i-1}(a)$. Recall that the capacities of all applicants in I^i are as in I^{i-1} except for the capacity of a^i that has been increased by 1. Hence, for all $a \in A \setminus \{a^i\}$, $|\xi(a)|$ does not exceed the capacity of a in instance I^{i-1} , namely $b^{i-1}(a)$, while $|\xi(a^i)|$ may exceed $b^{i-1}(a^i)$ by at most 1.

Moreover, it holds that $|\xi(a^i)| \geq |\mu^i(a^i)|$. Assuming to the contrary that $|\xi(a^i)| < |\mu^i(a^i)|$ yields that ξ is feasible also in instance I^{i-1} . In addition, $|\xi(a^i)| < |\mu^i(a^i)|$ implies that it cannot be the case that $\xi(a^i) \simeq_{a^i} \mu^i(a^i)$, and this, together with $\xi(a^i) \succeq_{a^i} \mu^i(a^i) \succeq_{a^i} \mu^{i-1}(a^i)$, yields $\xi(a^i) \succ_{a^i} \mu^i(a^i) \succeq_{a^i} \mu^{i-1}(a^i)$. But then ξ dominates μ^{i-1} in I^{i-1} , a contradiction to μ^{i-1} being a POM in I^{i-1} .

Let us first examine the case in which GSDT enters the ‘while’ loop and finds an augmenting path, hence μ^i dominates μ^{i-1} in I^i only with respect to applicant a^i that receives an additional course. This is one of her worst courses in $\mu^i(a^i)$ denoted as c_μ . Let c_ξ be a worst course for a^i in $\xi(a^i)$. Let also ξ' and μ' denote $\xi \setminus \{(a^i, c_\xi)\}$ and $\mu^i \setminus \{(a^i, c_\mu)\}$, respectively. Observe that both ξ' and μ' are feasible in I^{i-1} , while having shown that $|\xi(a^i)| \geq |\mu^i(a^i)|$ implies through Lemma 6 that ξ' weakly dominates μ' which in turn weakly dominates μ^{i-1} by Lemma 5. Since μ^{i-1} is a POM in I^{i-1} , $\xi'(a) \simeq_a \mu'(a) \simeq_a \mu^{i-1}(a)$ for all $a \in A$, therefore ξ dominates μ^i only with respect to a^i and $c_\xi \succ_{a^i} c_\mu$. Overall, $\xi(a) \simeq_a \mu^i(a) \simeq_a \mu^{i-1}(a)$ for all $a \in A \setminus \{a^i\}$ and $\xi(a^i) \succ_{a^i} \mu^i(a^i) \succ_{a^i} \mu^{i-1}(a^i)$.

Let t_ξ and t_μ be the ties of applicant a^i containing c_ξ and c_μ , respectively, where $t_\xi < t_\mu$ because $c_\xi \succ_{a^i} c_\mu$. Then, Lemma 5 implies that there is a path augmenting f^{i-1} (i.e., the flow corresponding to μ^{i-1}) in the network N^{i,t_ξ} . Let also t' be the value of $\text{curr}(a^i)$ at the beginning of stage i . Since we examine the case where GSDT enters the ‘while’ loop and finds an augmenting path, $C_{t'}^{a^i} \neq \emptyset$. Thus, t' indexes the least preferred tie from which a^i has a course in μ^{i-1} . The same holds for ξ' since $\xi'(a^i) \simeq_{a^i} \mu^{i-1}(a^i)$. Because ξ' is obtained by removing from a^i its worst course in $\xi(a^i)$, that course must belong to a tie of index no smaller than t' , i.e., $t' \leq t_\xi$. This together with $t_\xi < t_\mu$ yield $t' \leq t_\xi < t_\mu$, which implies that GSDT should have obtained ξ instead of μ^i at stage i , a contradiction.

It remains to examine the cases where, at stage i , GSDT does not enter the ‘while’ loop or enters it but finds no augmenting path. For both these cases, $\mu^i = \mu^{i-1}$, thus ξ dominating μ^i means that ξ is not feasible in I^{i-1} (since it would then also dominate μ^{i-1}). Then, it holds that $|\xi(a^i)|$ exceeds $b^{i-1}(a^i)$ by 1, thus $|\xi(a^i)| > |\mu^i(a^i)|$ yielding $\xi(a^i) \succ_{a^i} \mu^i(a^i)$. Let t_ξ be defined as above and t' now be the most preferred tie from which a^i has more courses in ξ than in μ^i . Clearly, $t' \leq t_\xi$ since t_ξ indexes the least preferred tie from which a^i has a course in ξ . If $t' < t_\xi$, then the matching ξ' , defined as above, is feasible in I^{i-1} and dominates μ^{i-1} because $\xi'(a^i) \succ_{a^i} \mu^{i-1}(a^i)$, a contradiction; the same holds if $t' = t_\xi$ and a^i has in ξ

at least two more courses from t_ξ than in μ^i . Otherwise, $t' = t_\xi$ and a^i has in ξ exactly one more course from t_ξ than in μ^i ; that, together with $|\xi(a^i)| > |\mu^i(a^i)|$ and the definition of t_ξ , implies that the index of the least preferred tie from which a^i has a course in μ^{i-1} and, therefore, the value of $\text{curr}(a^i)$ in the beginning of stage i , is at most t' . But then GSDT should have obtained ξ instead of μ^i at stage i , a contradiction. $\square$

The following statement is now direct.

Corollary 8 *GSDT produces a POM for instance I .*

To derive the complexity bound for GSDT, let us denote by L the length of the preference profile in I , i.e., the total number of courses in the preference lists of all applicants. Notice that $|E_3| = L$ and neither the size of any matching in I nor the total number of ties in all preference lists exceeds L .

Within one stage, several searches in the network might be needed to find a tie of the active applicant for which the current flow can be augmented. However, one tie is unsuccessfully explored at most once, hence each search either augments the flow thus adding a pair to the current matching or moves to the next tie. So the total number of searches performed by the algorithm is bounded by the size of the obtained matching plus the number of ties in the preference profile, i.e., it is $O(L)$. A search requires a number of steps that remains linear in the number of arcs in the current network (i.e., $N^{i, \text{curr}(a^i)}$), but as at most one arc per E_1 , E_2 and E_4 is used, any search needs $O(|E_3|) = O(L)$ steps. This leads to a complexity bound $O(L^2)$ for GSDT.

The next theorem will come in handy when implementing Algorithm 1, as it implies that for each applicant a only one node in T corresponding to a has to be maintained at a time.

Theorem 9 *Let $N^{i,t}$ be the network at stage i while tie t of applicant a^i is examined. Then, there is no augmenting path with respect to f^{i-1} in $N^{i,t}$ that has an arc of the form $((a^j, \ell), c)$ where $j \leq i$ and $\ell < \text{curr}(a^j)$.*

Proof Assume otherwise. Let $\mathfrak{P}$ be such an augmenting path that is found in round i and used to obtain μ^i . Hence $\mathfrak{P}$ corresponds to an augmenting path coalition $\mathfrak{C}$ of the following form $\langle a^i, c_r, \dots, c_s, a^j, c_q, \dots \rangle$ where c_r is in the t' th indifference class of a^i and both c_s and c_q are in the ℓ 'th indifference class of a^j , $\ell < \text{curr}(a^j)$. Note that as $\ell \geq 1$ thus $\text{curr}(a^j) > 1$. It then follows from the description of Algorithm 1 that either a^j is matched to at least one course in $\text{curr}(a^j)$ under μ^i or she is exposed (which would be the case when $\text{curr}(a^j) > n_2$ or $C_{\text{curr}(a^j)}^{a^j} = \emptyset$).

We first show that a^i and a^j are not the same applicant. Otherwise, $\mathfrak{C}' = \langle a^j, c_q, \dots \rangle$ —obtained from $\mathfrak{C}$ by discarding all courses and applicants that appear before a^j —is an augmenting path coalition with respect to μ^{i-1} in I^i . Clearly the matching obtained from μ^{i-1} by satisfying $\mathfrak{C}'$ Pareto dominates μ^i , as $\ell < t$, contradicting that μ^i is a POM in I^i .

In the remainder of the proof we assume that $a^i \neq a^j$. Let us first consider the case where a^j is matched to a course in $C_{curr(a^j)}^{a^j}$ under μ^i , and therefore by Lemma 1 to a course $c \in C_{curr(a^j)}^{a^j}$ under μ^{i-1} . Let $\mathcal{C}' = \langle c, a^j, c_q, \dots \rangle$, i.e., $\mathcal{C}'$ is obtained from $\mathcal{C}$ by discarding all courses and applicants that appear before a^j and replacing them with c . Since $\ell < curr(a^j)$, we have $c_q \succ_{a^j} c$. It is then easy to see that $\mathcal{C}'$ is an alternating path coalition with respect to μ^{i-1} in I^{i-1} , and hence μ^{i-1} is not a POM in I^{i-1} , a contradiction. We now consider the case where a^j is exposed in μ^i , and hence exposed in μ^{i-1} with respect to I^{i-1} . Let $\mathcal{C}' = \langle a^j, c_q, \dots \rangle$, i.e., $\mathcal{C}'$ is obtained from $\mathcal{C}$ by discarding all courses and applicants that appear before a^j . It is clear that $\mathcal{C}'$ is an augmenting path coalition with respect to μ^{i-1} in I^{i-1} , contradicting that μ^i is a POM in I^{i-1} . $\square$

Next we show that GSDT can produce any POM. Our proof makes use of a subgraph of the extended envy graph of Definition 1.

Theorem 10 *Given a CA instance I and a POM μ , there exists a suitable priority ordering over applicants Σ given which GSDT can produce μ .*

Proof Given an instance I and a POM μ , let $G = (V, E)$ be a digraph such that $V = \{ac : (a, c) \in \mu\}$ and there is an arc from ac to $a'c'$ if $a \neq a'$, $c' \notin \mu(a)$ and $c' \succeq_a c$. An arc $(ac, a'c')$ has weight -1 if a prefers c' to c and has weight 0 if she is indifferent between the two courses. Note that G is a subgraph of the extended envy graph $G(\mu)$ introduced in Definition 1. We say that ac envies $a'c'$ if $(ac, a'c')$ has weight -1 .

Note that if there exists an applicant-course pair $ac \in \mu$ and a course c' such that $c' \succ_{ac}$ and $ac' \notin \mu$, then c' must be full under μ , or else μ admits an alternating path coalition and is not a POM.

Moreover, all arcs in any given strongly connected component (SCC) of G have weight 0 . To see this, note that by the definition of a SCC, there is a path from any node to every other node in the component. Hence, if there is an arc in a SCC with weight -1 , then there must be a cycle of negative weight in that SCC. It is then straightforward to see that, by Lemma 1, μ admits a cyclic coalition, a contradiction.

It then follows that if $c' \succ_{ac}$, then ac' and ac cannot belong to the same SCC. Should that occur, there would be an arc (of weight 0) from ac' to some vertex $a'c^*$ in the same SCC, implying that $c^* \succeq_a c'$ and thus $c^* \succ_{ac}$. This in turn would yield that there is an arc of weight -1 from ac to $a'c^*$ in this SCC, a contradiction.

We create the condensation G' of the graph G . It follows from the definition of SCCs that G' is a DAG. Hence G' admits a topological ordering. Let X' be a reversed topological ordering of G' and X be an ordering of the vertices in G that is consistent with X' (the vertices within one SCC may be ordered arbitrarily). Let Σ be an ordering over applicant copies that is consistent with X (we can think of it as obtained from X by removing the courses). We show that GSDT can produce μ given Σ . Note that

technically Σ must contain $b(a)$ copies of each applicant a . However, as μ is Pareto optimal, upon obtaining μ the algorithm will not be able to allocate any more courses to any of the applicants, even if we append more applicant copies to the end of Σ .

We continue the proof with an induction. Recall that X is an ordering over the matched pairs in μ . We let $X(a^i) = c$ where $a^i c$ is the i th element of X . Let μ^i denote the matching that corresponds to the first i elements of X ; hence $\mu^{|\mu|} = \mu$. We claim that given Σ , GSDT is able to produce μ^i at stage i , after augmenting f^{i-1} through an appropriate augmenting path.

For the base case, note that $a^1 X(a^1)$ does not envy any other vertex in G and hence it can only be that $X(a^1) \in C_{a^1}^1$. It is then easy to see that the path $\langle \sigma, a^1, (a^1, 1), X(a^1), \tau \rangle$ is a valid augmenting path in $N^{1,1}$ and hence GSDT might choose it.

Assume that GSDT produces μ^ℓ at the end of each stage ℓ for all $\ell < i$. We prove that it can produce μ^i at the end of stage i . Assume, for a contradiction, that this is not the case. Let r denote the indifference class of a^i to which $X(a^i)$ belongs. Note that course $X(a^i)$ is not full in I^i , so if the path $\langle \sigma, a^i, (a^i, r), X(a^i), \tau \rangle$ is not chosen by GSDT, it must be the case that GSDT finds an augmenting path in $N^{i,t}$ for some $t < r$. Let $\mathfrak{C}$ denote the corresponding augmenting path coalition with respect to the matching μ^{i-1} , which would be of the form

$$a^i, c_{j_1}, a^{i_1}, \dots, c_{i_{y-1}}, a^{i_{y-1}}, c_{j_y}$$

where $c_{j_1} \in C_{a^i}^t$ and c_{j_y} is exposed in μ^{i-1} . It follows, from the definition of an augmenting path, that there is an edge $(a^{i_k} c_{j_k}, a^{i_{k+1}} c_{j_{k+1}})$ in G for all k , $1 \leq k \leq y-2$. Furthermore, as $c_{j_1} \succ_{a^i} \mu(a^i)$, there is an edge of weight -1 in G from $a^i X(a^i)$ to $a^{i_1} c_{j_1}$; therefore $a^{i_1} c_{j_1}$ belongs to a SCC of higher priority than the one to which $a^i X(a^i)$ belongs. If c_{j_y} is exposed in μ , then $\mathfrak{C}'$ that is obtained by adding $X(a^i)$ to the beginning of $\mathfrak{C}$ is an alternating path coalition in μ , a contradiction to μ being a POM. Therefore c_{j_y} must be full in μ .

As c_{j_y} is not full in μ^{i-1} and full in μ , there must exist an a^z , $z > i$, such that $(a^z, c_{j_y}) \in \mu$. It follows, from the augmenting path coalition $\mathfrak{C}$, that there is a path in G from $a^{i_1} c_{j_1}$ to $a^z c_{j_y}$. If there is also a path from $a^z c_{j_y}$ to $a^{i_1} c_{j_1}$, then the two vertices belong to the same SCC; as $a^{i_1} c_{j_1}$ belongs to a SCC of higher priority than the one to which $a^i X(a^i)$ belongs, so must $a^z c_{j_y}$, implying that $a^z c_{j_y}$ must have appeared before $a^i X(a^i)$ in X , a contradiction to $z > i$. If there is no such a path, then $a^z c_{j_y}$ belongs to a SCC that is prioritized even over the SCC to which $a^{i_1} c_{j_1}$ belongs, and hence must have appeared before $a^i X(a^i)$ in X , a contradiction to $z > i$. $\square$

5 Truthfulness of Mechanisms for sPOMs

It is well-known that the SDM for HA is truthful, regardless of the given priority ordering over applicants. We will show shortly that GSDT is not necessarily

truthful, but first prove that this property does hold for some priority orderings over applicants.

Theorem 11 *GSDT is truthful given Σ if, for each applicant a , all occurrences of a in Σ are consecutive.*

Proof Without loss of generality, let the applicants appear in Σ in the following order:

$$\underbrace{a_1, a_1, \dots, a_1}_{b(a_1)\text{-times}}, \underbrace{a_2, a_2, \dots, a_2}_{b(a_2)\text{-times}}, \dots, \underbrace{a_{i-1}, a_{i-1}, \dots, a_{i-1}}_{b(a_{i-1})\text{-times}}, \underbrace{a_i, a_i, \dots, a_i}_{b(a_i)\text{-times}}, \dots$$

Assume to the contrary that some applicant benefits from misrepresenting her preferences. Let a_i be the first such applicant in Σ who reports $P'(a_i)$ instead of $P(a_i)$ in order to benefit and $\mathcal{P}' = (P'(a_i), \mathcal{P}(-a_i))$. Let μ denote the matching returned by GSDT using ordering Σ on instance $I = (A, C, \mathcal{P}, b, q)$ (i.e. the instance in which applicant a_i reports truthfully) and ξ the matching returned by GSDT using Σ but on instance $I' = (A, C, \mathcal{P}', b, q)$. Let $s = (\sum_{\ell < i} b(a_\ell)) + 1$, i.e., s is the first stage in which our mechanism considers applicant a_i . Let j be the first stage of GSDT such that a_i prefers ξ^j to μ^j , where $s \leq j < s + b(a_i)$.

Given that applicants $a_1, \dots, a_{i-1}$ report the same in I as in I' and all their occurrences in Σ are before stage j , Lemma 5 yields $\mu^j(a_\ell) \simeq_{a_\ell} \xi^j(a_\ell)$ for $\ell = 1, 2, \dots, i - 1$. Also $\mu^j(a_\ell) = \xi^j(a_\ell) = \emptyset$ for $\ell = i + 1, i + 2, \dots, n_1$, since no such applicant has been considered before stage j . But then, all applicants apart from a_i are indifferent between μ^j and ξ^j , therefore a_i preferring ξ^j to μ^j implies that μ^j is not a POM in I^j , a contradiction to Theorem 7. $\square$

The next result then follows directly from Theorem 11.

Corollary 12 *GSDT is truthful if all applicants have quota equal to one.*

There are priority orderings for which an applicant may benefit from misreporting her preferences, even if preferences are strict. This phenomenon has also been observed in a slightly different context [6]. Let us also provide an example.

Example 13 Consider a setting with applicants a_1 and a_2 and courses c_1 and c_2 , for which $b(a_1) = 2$, $b(a_2) = 1$, $q(c_1) = 1$, and $q(c_2) = 1$. Let I be an instance in which $c_2 \succ_{a_1} c_1$ and a_2 finds only c_1 acceptable. This setting admits two POMs, namely $\mu_1 = \{(a_1, c_2), (a_2, c_1)\}$ and $\mu_2 = \{(a_1, c_1), (a_1, c_2)\}$.

GSDT returns μ_1 for $\Sigma = (a_1, a_2, a_1)$. If a_1 misreports by stating that she prefers c_1 to c_2 , GSDT returns μ_2 instead of μ_1 . Since $\mu_2 \succ_{a_1} \mu_1$, GSDT is not truthful given Σ .

The above observation seems to be a deficiency of GSDT. We conclude by showing that no mechanism capable of producing all POMs is immune to this shortcoming.

Fig. 2 Four instances of CA used in the proof of Theorem 14. In all four instances $b(a_1) = 2$, $b(a_2) = 1$, $q(c_1) = q(c_2) = 1$. For each of instances I_1 to I_3 , a matching is indicated using circles in applicants' preference lists

Theorem 14 *There is no universally truthful randomized mechanism that produces all POMs in CA, even if applicants' preferences are strict and all courses have quota equal to one.*

Proof The instance I_1 in Fig. 2 admits three POMs, namely $\mu_1 = \{(a_1, c_1), (a_2, c_2)\}$, $\mu_2 = \{(a_1, c_1), (a_1, c_2)\}$ and $\mu_3 = \{(a_1, c_2), (a_2, c_1)\}$. Assume a randomized mechanism ϕ that produces all these matchings. Therefore, there must be a deterministic realization of it, denoted as ϕ^D , that returns μ_1 given I_1 . Let us examine the outcome of ϕ^D under the slightly different applicants' preferences shown in Fig. 2, bearing in mind that ϕ^D is truthful.

- Under I_2 , ϕ^D must return μ_2 . The only other POM under I_2 is μ_3 , but if ϕ^D returns μ_3 , then a_2 under I_1 has an incentive to lie and declare only c_1 acceptable (as in I_2).
- Under I_3 , ϕ^D must return μ_2 . The only other POM under I_3 is μ_3 , but if ϕ^D returns μ_3 , then a_1 under I_3 has an incentive to lie and declare that she prefers c_1 to c_2 (as in I_2).

I_4 admits two POMs, namely μ_2 and μ_3 . If ϕ^D returns μ_2 , then a_1 under I_1 has an incentive to lie and declare that she prefers c_2 to c_1 (as in I_4). If ϕ^D returns μ_3 , then a_2 under I_3 has an incentive to lie and declare c_2 acceptable—in addition to c_1 —and less preferred than c_1 (as in I_4). Thus overall ϕ^D cannot return a POM under I_4 while maintaining truthfulness. $\square$

6 Future Work

Our work raises several questions. A particularly important problem is to investigate the expected size of the matching produced by the randomized version of GSdT. It is also interesting to characterize priority orderings that imply truthfulness in GSdT. Consequently, it will be interesting to compute the expected size of the matching produced by a randomized GSdT in which the randomization is taken over the priority orderings that induce truthfulness.

Acknowledgments This research has been co-financed by the European Union (European Social Fund - ESF) and Greek national funds under Thales grant MIS 380232 (Eirinakis, Mourtos), by grant EP/K010042/1 from the Engineering and Physical Sciences Research Council (Manlove, Rastegari), grants VEGA 1/0344/14, 1/0142/15 from the Slovak Scientific grant agency VEGA (Cechlárová), student grant VVGS-PF-2014-463 (Oceláková) and OTKA grant K108383 (Fleiner). The authors gratefully acknowledge the support of COST Action IC1205 on Computational Social Choice.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Abdulkadiroğlu, A., Sönmez, T.: Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* **66**(3), 689–701 (1998)
2. Abraham, D.J., Cechlárová, K., Manlove, D.F., Mehlhorn, K.: Pareto optimality in house allocation problems. In: Proc. ISAAC '04, volume 3341 of LNCS, pp. 3–15. Springer (2004)
3. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., Upper Saddle River (1993)
4. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. *J. Econ. Theory* **100**(2), 295–328 (2001)
5. Bogomolnaia, A., Moulin, H.: Random matching under dichotomous preferences. *Econometrica* **72**(1), 257–279 (2004)
6. Budish, E., Cantillon, E.: The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *Amer. Econ. Rev.* **102**(5), 2237–71 (2012)
7. Cechlárová, K., Eirínakis, P., Fleiner, T., Magos, D., Manlove, D., Mourtos, I., Oceláková, E., Rastegari, B.: Pareto optimal matchings in many-to-many markets with ties. In: Proceedings of the Eighth International Symposium on Algorithmic Game Theory (SAGT'15), pp. 27–39 (2015)
8. Cechlárová, K., Eirínakis, P., Fleiner, T., Magos, D., Mourtos, I., Potpinková, E.: Pareto optimality in many-to-many matching problems. *Discret. Optim.* **14**(0), 160–169 (2014)
9. Chen, N., Ghosh, A.: Algorithms for Pareto stable assignment. In: Conitzer, V., Rothe, J. (eds.) Proc. COMSOC '10, pp. 343–354. Düsseldorf University Press (2010)
10. Cherkassky, B.V., Goldberg, A.V.: Negative-cycle detection algorithms. *Math. Program.* **85**, 227–311 (1999)
11. Fujita, E., Lesca, J., Sonoda, A., Todo, T., Yokoo, M.: A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences. In: Proc. AAAI '15 (2015)
12. Gärdenfors, P.: Assignment problem based on ordinal preferences. *Manag. Sci.* **20**(3), 331–340 (1973)
13. Gigerenzer, G., Goldstein, D.G.: Reasoning the fast and frugal way: Models of bounded rationality. *Psychol. Rev.* **103**(4), 650–669 (1996)
14. Goldberg, A.V.: Scaling algorithms for the shortest paths problem. Technical Report STAN-CS-92-1429, Stanford University (Stanford, CA, US) (1992)
15. Hylland, A., Zeckhauser, R.: The efficient allocation of individuals to positions. *J. Polit. Econ.* **87**(2), 293–314 (1979)
16. Klaus, B., Miyagawa, E.: Strategy-proofness, solidarity, and consistency for multiple assignment problems. *Int. J. Game Theory* **30**, 421–435 (2001)
17. Krysta, P., Manlove, D., Rastegari, B., Zhang, J.: Size versus truthfulness in the House Allocation problem. Technical Report 1404.5245, Computing Research Repository, Cornell University Library. A shorter version appeared in the Proceedings of EC'14 (2014)
18. Manlove, D.F.: *Algorithmics of Matching Under Preferences*. World Scientific (2013)
19. Saban, D., Sethuraman, J.: The complexity of computing the random priority allocation matrix. In: Proc. WINE '13, volume 8289 of LNCS, pp. 421. Springer (2013)
20. Svensson, L.G.: Queue allocation of indivisible goods. *Soc. Choice Welfare* **11**(4), 323–330 (1994)
21. Zhou, L.: On a conjecture by Gale about one-sided matching problems. *J. Econ/ Theory* **52**(1), 123–135 (1990)