Repeatedly Matching Items to Agents Fairly and Efficiently

Extended Abstract

Ioannis Caragiannis
Aarhus University
Aarhus, Denmark
iannis@cs.au.dk

Shivika Narang
Indian Institute of Science
Bengaluru, India
shivika@iisc.ac.in

ABSTRACT

We consider a novel setting where a set of items are matched to the same set of agents repeatedly over multiple rounds. Each agent gets exactly one item per round, which brings interesting challenges to finding efficient and/or fair *repeated matchings*. A particular feature of our model is that the value of an agent for an item in some round depends on how often the item has been used by the agent in the past. We present a set of positive and negative results about the efficiency and fairness of repeated matchings.

KEYWORDS

Resource allocation; fair division; envy-freeness; matchings

ACM Reference Format:
Ioannis Caragiannis and Shivika Narang. 2023. Repeatedly Matching Items to Agents Fairly and Efficiently: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The problem of fairly dividing indivisible items among agents has received enormous attention by the EconCS research community in recent years[1–4, 7, 8]. The standard setting involves a set of items and agents who have values for them. The objective is to compute a fair allocation which gives each item to a single agent. In practice, sometimes, the same set of items must be allocated to the same set of agents repeatedly. Prior work has typically explored various settings where agents' allocations do not change with time, with few exceptions [6, 7]. More crucially, another feature that distinguishes such scenarios from the standard setting is that the value of an agent for an item changes over time and typically depends on how many times the agent has received the item in the past.

To give an example, consider different research labs that all need access to several expensive research facilities in a university. The access to these facilities must be fairly coordinated/scheduled throughout the year. To be fair among labs and efficient overall, we must take into account the values the labs have for facilities, which typically change over time. For instance, during the first few weeks of access to a facility, the researchers in a lab may need time to learn how to operate it. During that time, the value the lab gets by accessing a facility can be very low, even negative. As the researchers gain more experience, their research output increases, and so does the lab's value for the facility. Once the researchers have run their intended experiments, the lab's value for the facility decreases again until the next experiment.

Conceptual contribution. To capture such situations, we introduce a new model of *repeated matchings* with n agents who must be matched with exactly one of n items in each of T rounds, repeatedly. An important novelty of our model is that valuations are *history-dependent*: the value an agent has for an item in a round depends on how many times the agent has used the item in previous rounds. We use *social welfare* to assess the efficiency of repeated matchings. We also use relaxations of *envy-freeness* as fairness concepts. We adapt the well-known *envy-freeness up to one item* (EF1) and use it when all valuations are non-negative (i.e., when items are *goods*). We observe that EF1 is not suitable when valuations can be positive or negative (i.e., when items are *mixed*), and introduce the notion of *swap envy-freeness* to assess fairness of repeated matchings for mixed items.

Technical contribution. We prove that the problem of computing a repeated matching with maximum social welfare is NP-hard, even when $T = 3$. Our hardness reduction defines instances with non-monotone valuations. The problem becomes solvable in polynomial time when the valuations are monotone. We also consider fair repeated matchings. First, for EF1, we find that under identical valuations, EF1 repeated matchings always exist and can be found tractably. Furthermore, we show that any instance with general valuations and $T \bmod n \in \{0, 1, 2, n-1\}$ (i.e., including all instances with at most four agents/items) has an EF1 repeated matching, which can be computed efficiently. Unfortunately, EF1 is not compatible with social welfare maximization and even approximating the maximum social welfare over EF1 repeated matchings is NP-hard. This holds even when EF1 solutions and social welfare maximizing solutions can be found in polynomial time separately.

We further propose and study a new fairness notion called swap envy-freeness (swapEF). Here, we find that under identical valuations, swapEF repeated matchings can be found using the same algorithm as used for EF1. Furthermore, we show that swapEF repeated matchings always exist and can be computed efficiently on instances with $T \bmod n \in \{0, 1, 2, n-2, n-1\}$ (i.e., including all instances with at most five agents/items). Our hardness results are proved on instances with goods. Our positive results besides those for EF1, apply to instances with mixed items. The complete details can be found in [5].

2 NOTATION AND PRELIMINARIES

Our setting involves a set $\mathcal{A}$ of n agents and a set $\mathcal{G}$ of n items. We consider instances of the form $I = \langle \mathcal{A}, \mathcal{G}, \{v_i\}_{i\in\mathcal{A}}, T\rangle$, where T denotes the number of rounds and, for each agent $i \in \mathcal{A}$, v_i is a function from $\mathcal{G} \times [T]$ to $\mathbb{R}$, where $v_i(g, t)$ denotes the *valuation* of

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

agent i for item g when it is matched to the item for the t^{th} time. A repeated matching $A = (A^1, ..., A^T)$ is simply a collection of matchings, with one matching A^t per each round $t \in [T]$. Furthermore, we denote by A_i the multiset (or *bundle*) which contains copies of the items to which agent $i \in \mathcal{A}$ is matched in the T rounds.

Hence, defining the bundles A_i for $i \in \mathcal{A}$ given the repeated matching A is trivial. The opposite task is also straightforward. Let $N(B, g)$ be the multiplicity of item g in bundle B. Given bundles of items A_i for $i \in \mathcal{A}$ with $|A_i| = T$ (i.e., each agent gets T copies of items) and $\sum_{i \in \mathcal{A}} N(A_i, g) = T$ (i.e., T copies of each item g are allocated), a consistent repeated matching[1] for instance I is obtained as follows. We construct the bipartite multigraph $G = (\mathcal{A}, \mathcal{G}, E)$ so that the set of edges E consists of (a copy of) edge (i, g) for every (copy of) item g such that $g \in A_i$. The graph G is T-regular and, thus, by Hall's matching theorem (see 9), can be decomposed into T matchings of edges $M_1, ..., M_T$. These matchings correspond to a repeated matching by interpreting the edge (i, g) in matching M_t as the assignment of item g to agent i in the t^{th} round.

With a slight abuse of notation, we use $v_i(B)$ to denote the *value* agent $i \in \mathcal{A}$ has when she gets the bundle B, i.e., $v_i(B) = \sum_{g \in \mathcal{G}} \sum_{t=1}^{N(B,g)} v_i(g, t)$. Hence, for a repeated matching A, $v_i(A_i)$ is the total value from each item copy agent i receives in all rounds. The *social welfare* of A is simply the sum of the agents' values for their bundle, i.e., $SW(A) = \sum_{i \in \mathcal{A}} v_i(A_i)$.

We shall look at specific types of valuations under which we will try to find efficient and/or fair repeated matchings. A well-motivated setting is that of *identical* valuations where $v_1 = v_2 = \cdots = v_n$. This assumption proves particularly useful in finding fair solutions. Another important class of valuation functions is that of *monotone* valuations. We extend to repeated matchings *envy-freeness of up to one item* (EF1) as follows.

*Definition 2.1 (*EF1*).* A repeated matching A is EF1 if for every pair of agents $i, j \in \mathcal{A}$, there exists an item $g \in \mathcal{G}$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

We remark that the operation $A_j \setminus \{g\}$ removes one copy of item g from the bundle A_j if g belongs to A_j and leaves A_j intact otherwise.

We refer to the items as *goods* on instances where all valuations are non-negative, i.e., when $v_i(g, t) \geq 0$ for every $i \in \mathcal{A}$, $g \in \mathcal{G}$, and $t \in [T]$. When there are no restrictions on the valuations, we refer to the items as *mixed*. Consider the following instance with $n = 2$ and $T = 1$. One of the items is a good and the other is a chore. There are exactly two possible matchings. In either, the classical extension of EF1 for mixed items from the fair division literature (e.g., see 1), which requires that the value of an agent is higher than that of another either by removing a single item from either one of the two bundles, is not satisfied. Motivated by this simple example, we propose and investigate an alternate notion of fairness, which we call *swap envy-freeness* (swapEF).

*Definition 2.2 (*swapEF*).* Let $I = \langle \mathcal{A}, \mathcal{G}, \{v_i\}_{i \in \mathcal{A}}, T \rangle$ be a repeated matching instance with mixed items. A repeated matching $A = (A_1, ..., A_n)$ in I is swapEF if for every pair of agents $i, j \in \mathcal{A}$, either (i) or (ii) is true:

(i) $v_i(A_i) \geq v_i(A_j)$;
(ii) There exist items $g_i \in A_i$ and $g_j \in A_j$ such that $v_i(A_i \cup \{g_j\} \setminus \{g_i\}) \geq v_i(A_j \cup \{g_i\} \setminus \{g_j\})$.

Condition (ii) requires that the value agent i has for her bundle A_i after replacing a copy of item g_i with an extra copy of item g_j is at least as high as her value for the bundle A_j of agent j after exchanging a copy of item g_j with a copy of item g_i.

[1] We remark that this repeated matching is not unique. However, this does not affect the values of each agent for her bundle and the bundle of any other agent, which are the same in all different consistent repeated matchings.

3 FAIRNESS UNDER IDENTICAL VALUATIONS

Our algorithm starts by assigning $\lfloor T/n \rfloor$ copies of each item to each agent. If $T \bmod n > 0$ (i.e., additional copies have to be assigned to the agents so that the repeated matching is correct), the algorithm works in a round robin fashion for $T \bmod n$ phases. In these phases, it uses a fixed ranking of the items according to the value $v(g, \lceil T/n \rceil)$ of their $\lceil T/n \rceil$-th copy. The ranking assigns to each item a distinct integer $\text{rank}(g)$ in $[n]$ such that $\text{rank}(g_1) < \text{rank}(g_2)$ implies that $v(g_1, \lceil T/n \rceil) \geq v(g_2, \lceil T/n \rceil)$. In each round-robin phase, the agents act according to the ordering $1, 2, ..., n$. When it is agent i's turn, she picks a copy of the lowest-rank item that is available. The algorithm appears below as Algorithm 1. It has access to function rank() defined as above and uses the matrix f to store the number of copies of each item an agent gets. The final step is to call routine GenerateFromFreq() to transform f to the repeated matching A.

Algorithm 1: Fairness under identical valuations

Input: Identical Valuations Instance $I = \langle \mathcal{A}, \mathcal{G}, v, T \rangle$ with $|\mathcal{A}| = n$
Output: A repeated matching A
$f(i, g) \leftarrow \lfloor T/n \rfloor$, $\forall i \in \mathcal{A}$, $\forall g \in \mathcal{G}$;
if $T \bmod n > 0$ **then**
  $x_g \leftarrow T \bmod n$, $\forall g \in G$;
  for $t = 1$ *to* $T \bmod n$ **do**
    for $i = 1$ *to* n **do**
      $g' \leftarrow \arg\min_{g: x_g > 0} \text{rank}(g)$;
      $x_{g'} \leftarrow x_{g'} - 1$;
      $f(i, g') \leftarrow \lceil T/n \rceil$;
$A \leftarrow \text{GenerateFromFreq}(f)$;

We can use Algorithm 1 to prove the next statement.

Theorem 3.1. *Given a repeated matching instance with identical valuations, for goods, an* EF1 *repeated matching exists and can be computed in polynomial time. For mixed items, a* swapEF *repeated matching exists and can be coputed in polynomial time.*

The proof along with our other contributions can be found in [5].

4 OPEN PROBLEMS

Our work leaves several interesting open problems that deserve investigation. Understanding social welfare maximization is the first one. Is the problem hard for instances with two rounds? The problem is in P for a single round and NP-hard for $T = 3$. What about approximation algorithms when the items are goods and valuations are not necessarily monotone? Is a constant approximation ratio possible? Regarding fairness, the most important open question is whether EF1 repeated matchings exist for any instance with goods. Furthermore, are EF1 and swapEF compatible with different notions of efficiency than the utilitarian social welfare we have used here?

ACKNOWLEDGMENTS

Ioannis Caragiannis is partially supported by the Independent Research Fund Denmark (DFF) under grant 2032-00185B. Shivika Narang is supported by a research fellowship from Tata Consultancy Services.

REFERENCES

[1] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2022. Fair allocation of indivisible goods and chores. *Autonomous Agents & Multi Agent Systems* 36, 1 (2022), art. 3.

[2] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair division of indivisible goods among strategic agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1811–1813.

[3] Siddharth Barman, Anand Krishna, Pooja Kulkarni, and Shivika Narang. 2021. Sublinear approximation algorithm for Nash social welfare with XOS valuations. *arXiv preprint arXiv:2110.00767* (2021).

[4] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 1–32.

[5] Ioannis Caragiannis and Shivika Narang. 2022. Repeatedly matching items to agents fairly and efficiently. *arXiv preprint arXiv:2207.01589* (2022).

[6] Sreenivas Gollapudi, Kostas Kollias, and Benjamin Plaut. 2020. Almost envy-free repeated matching in two-sided markets. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*. 3–16.

[7] Hadi Hosseini, Kate Larson, and Robin Cohen. 2015. Matching with dynamic ordinal preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 936–943.

[8] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.

[9] Michael D. Plummer and László Lovász. 1986. *Matching Theory*. Elsevier Science.