

Fairness in Multiagent Resource Allocation with Dynamic and Partial Observations

Extended Abstract

Aurélie Beynier, Nicolas Maudet and Anastasia Damamme

Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France
{aurelie.beynier,nicolas.maudet,anastasia.damamme}@lip6.fr

ABSTRACT

We investigate fairness issues in distributed resource allocation of indivisible goods. More specifically, we study envy-freeness in a setting where the observations of agents only result from encounters with other agents. Agents thus have a partial and uncertain view of the entire allocation, that they maintain throughout the process, and which allows them to have different estimates of their envy. We provide a fully distributed protocol allowing to guarantee termination despite the limited knowledge of agents.

ACM Reference Format:

Aurélie Beynier, Nicolas Maudet and Anastasia Damamme. 2018. Fairness in Multiagent Resource Allocation with Dynamic and Partial Observations. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10-15, 2018*, IFAAMAS, 3 pages.

1 INTRODUCTION

Fairly allocating indivisible resources is an ubiquitous problem, with applications ranging from course allocation, divorce settlement, or cloud computing [3, 5, 12]. In many contexts though, agents may not perceive the entire system (or even if they could, they often don't have the cognitive or computational ability to do so). This motivated several recent works to reconsider the classical notions of fairness, like for instance *proportionality* or *envy*, in settings where the knowledge of agents is partial [1, 2, 6, 7]; for instance only the resources held by neighbours in a network can be observed.

In this work we depart from these previous studies by investigating distributed protocols in a setting where agents have local but also time-stamped views of the current allocation within the system. Indeed, we consider that the visibility graph of agents is not fixed but results instead from an endogenous dynamics of the system: agents meet each other (in a pairwise manner) to perform (mutually beneficial) local deals [9, 14]. When they do so they observe the current bundle held by the other agent. As a result, the view of agents is not only *incomplete*, but also *uncertain*, as an agent observations may become obsolete after a while. Our experimental results (not reported here for lack of space) compare different heuristics that agents may use to select their encounter, in particular in terms of the fairness of outcomes and the number of messages induced.

2 OUR MODEL

We consider a set $\mathcal{N} = \{1, \dots, n\}$ of agents and a set $\mathcal{R} = \{r_1, \dots, r_m\}$ of indivisible resources. Each agent holds initially a fixed number

of resources. To make the presentation easier, we shall take this number to be k for all agents (thus $m = k \cdot n$). However, our approach generalizes to agents holding different (but fixed throughout the process) number of resources.

An allocation A is a partition of the resources to the agents. A_i denotes the set of resources held by agent i in the allocation A . Each agent i has a utility function $u_i : \mathcal{R} \rightarrow \mathbb{R}^+$ mapping bundles of resources to reals. We adopt the commonly used framework of additive utility functions to represent preferences over the resources [11, 13]. The utility of the agent i for an allocation A is thus defined as the sum of the utilities over the resources forming A_i .

System dynamics. Agents start from an initial allocation, and modify it incrementally by exchanging resources [14]. Deals consist in (bilateral) swaps [8], by which two agents exchange one resource against another resource. We focus on rational agents which only accept to perform individually rational deals, albeit in a cooperative way [9]. The dynamics of our system is guided by the agents themselves, and is best described at two levels: (1) at the *global* level, agents decide to contact another agent; and (2) at the *local* level –once a bilateral contact has been established– the agents try to exchange resources. The system is *stable* when no more swap is possible among any agents. Convergence is easily guaranteed because the number of deals is finite, and each deal must strictly improve the utility of one of the agents [9, 14]. Hence, it would be sufficient to enforce at the global level a protocol which makes sure that each pairwise encounter is checked at some point. However, this approach is unsatisfactory in our setting. First, this contradicts the fact that our dynamics is endogenous. Secondly, this may be very inefficient in practice. An important issue is thus to design an efficient distributed technique allowing to detect convergence. Indeed, since the resource allocation process is distributed, we would like each agent to detect on his-own that the system has converged to a stable allocation.

Incomplete and incorrect knowledge. The observations of the agents are updated upon encounters: agents become aware of their respective bundles when they meet, but have no way of knowing how the allocation evolves beyond this. The principle that we shall use is the following: *unless proven otherwise, agents assume resources are still held where they were last observed*. Hence their views may not only be incomplete, but also *incorrect*.

Let us denote by O_i^j the up-to-date set of resources that i assigns to agent j . The update process is captured by two rules:

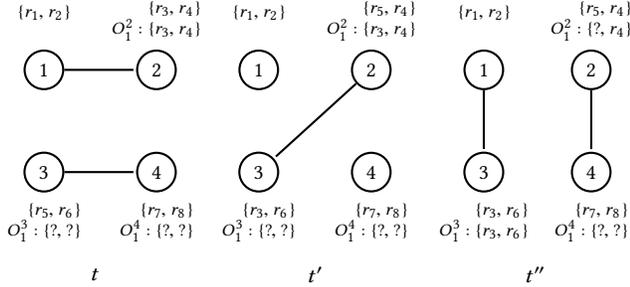
- (1) upon encountering an agent j , an agent i observes the k items $\{r_1, \dots, r_k\}$ that agent j currently holds, and thus updates $O_i^j \leftarrow \{r_1, \dots, r_k\}$

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10-15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

- (2) if an item r_l is observed by i upon encountering j while it was supposed to be held by j' ($r_l \in O_i^j$), then the observation set of i is updated: $O_i^j \leftarrow O_i^j \setminus \{r_l\}$

Now denote $\bar{O}_i = R \setminus \cup_{j \in N} O_i^j$, the set of items that agent i does not know where to allocate.

EXAMPLE 1. We picture a scenario involving four agents, and three time-steps ($t'' > t' > t$). We take the point of view of agent 1. Edges represent encounters between the agents.



At time t , agent 1 updates his observation set for agent 2. At time t' , an encounter takes place between agent 2 and agent 3, but agent 1 is not aware of this. Hence, out of the two items that agent 1 can assign to someone in his observation set, only one is correct. Finally, at time t'' , agent 1 encounters agent 3 and updates his observation sets for agent 3, but also for agent 2.

Fairness based on observations. We adapt the notion of *envy* [4, 10, 15] to uncertain and incomplete knowledge. The basic envy notion that we shall use is the pairwise degree of envy [11], but relative to a given set of observations, i.e. $e_{ij} = \max(0, u_i(O_i^j) - u_i(A))$

We call this notion *evidence-based envy*. A system is *evidence-based envy-free* (EBEF) when no agent is envious, based on his observations only. Because of the possible incorrectness of knowledge, it is easy to see that “actual” envy-freeness (as would be evaluated by an omniscient agent observing the true and complete allocation) does not imply evidence-based envy, nor vice-versa.

As agents hold a fixed number of resources, it is natural to consider different ways to complete the observations an agent may have regarding another agent in order to estimate the envy. Denote by $avg(O)$ the average value of items in O . Let $\bar{O}_i \uparrow [q]$ (resp. $\bar{O}_i \downarrow [q]$) be the top- q (resp. last- q) elements of \bar{O}_i , that is, the items not allocated with the q highest (resp. lowest) utility for agent i .

- *optimistic* envy of agent j , obtained by completing the missing items by the least valuable ones:

$$e_{ij}^{OPT} = \max(0, u_i(O_i^j \cup \bar{O}_i \downarrow [k - |O_i^j|]) - u_i(A))$$

- *pessimistic* envy of agent j , obtained by completing the missing items by the most valuable ones:

$$e_{ij}^{PES} = \max(0, u_i(O_i^j \cup \bar{O}_i \uparrow [k - |O_i^j|]) - u_i(A))$$

- *average* envy of agent j , obtained by completing the missing items by the average value of the set \bar{O}_i :

$$e_{ij}^{AV} = \max(0, (u_i(O_i^j) + (k - |O_i^j|) \cdot avg(\bar{O}_i)) - u_i(A))$$

Clearly, for any j , it is the case that $e_{ij}^{OPT} \leq e_{ij}^{AV} \leq e_{ij}^{PES}$, and all notions coincide with classical envy when the observation set of an agent is complete and correct.

EXAMPLE 2. Consider again Ex. 1 with the following preferences:

	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8
ag. 1	5	5	8	3	4	1	7	7

Let us consider agent 1. At time t , upon meeting agent 2, he becomes envious of agent 2 (as agent 2 holds $\{r_3, r_4\}$, a bundle he values 11). Now at time t' (recall that an exchange took place between agents 2 and 3 in the meantime), agent 1 meets agent 3. Agent 1 is not EBEF of agent 2. In fact, he is not even pessimistically envious of agent 2, since the value of items in the observation set for agent 2 is only 3, and the highest valued item in \bar{O}_1 ($\bar{O}_1 \uparrow [1]$) has utility 7. Interestingly, agent 1 is (even optimistically) envious of agent 4. In fact, agent 4 must either hold r_7 and r_8 together, or only one of these resources together with r_5 . To put it differently, $\bar{O}_1 = \{r_5, r_7, r_8\}$. As agent 4 must hold two resources, he holds a bundle of value at least 11.

3 DISTRIBUTED TERMINATION DETECTION

The agents try to agree on bilateral swap deals as long as some deals are still possible. Since each agent has limited observability of the system and does not know the preferences of the other agents, detecting the end of the exchanges in a distributed way is not easy. We propose a distributed approach for termination detection where each agent i maintains an *interest information set* about the resources of the system. In this set, each resource r_k is labeled as:

- *unattractive* (UN), meaning that resource r_k is not of interest i.e. $\forall r_l \in A_i, u_i(r_l) > u_i(r_k)$.

For the other (attractive) resources, two further labels are used:

- *to-try* (TT), meaning that the agent may try to obtain r_k ;
- *wait-for-new-resources* (WR), meaning that the agent has already tried to obtain r_k , and that the exchange failed. The agent must thus acquire new resources in order to propose (potentially) better exchanges to the agent holding r_k .

Initially, each agent distinguishes resources he may try to obtain -TT- from uninteresting resources -UN-. Each time an agent i encounters another agent j and no exchange of resources takes place, labels on the resources of the other agent are turned to WR. If the encounter has led to an exchange, some resources may now not be of interest for the agent. The information set is then re-initialized.

Each agent alternates between two execution modes: the *active mode* where the agent contacts the other agents and try to exchange resources and the *standby mode* where the agent waits for some contact requests from the other agents. Initially, each agent starts in the active mode. When an agent has no more resource labeled as TT, he moves to the standby mode. If another agent contacts him and at least one resource turns to TT, the agent comes out of the standby mode.

PROPOSITION 1. When all agents are in the standby mode, there is no more possible rational exchange of resources.

ACKNOWLEDGMENTS

This work is partially supported by the ANR project 14-CE24-0007-01 - CoCoRICo-CoDec.

REFERENCES

- [1] R. Abebe, J. Kleinberg, and D. Parkes. Fair Division via Social Comparison. In *Proc. 16th Intl. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2017)*, 2017.
- [2] H. Aziz, S. Bouveret, I. Caragiannis, I. Giagkousi, and J. Lang. Knowledge, fairness, and social constraints. In *Proceedings of the 32nd AAAI conference on Artificial Intelligence (AAAI'18)*, New Orleans, Louisiana, USA, 2018. AAAI Press. (to appear).
- [3] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016.
- [4] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, Mar. 2016.
- [5] S. Brams and A. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [6] Y. Chen and N. Shah. Ignorance is often bliss: Envy with incomplete information. Working paper, Harvard University, 2017.
- [7] Y. Chevaleyre, U. Endriss, and N. Maudet. Distributed Fair Allocation of Indivisible Goods. *Artificial Intelligence*, 242:1–22, Jan. 2017.
- [8] A. Damamme, A. Beynier, Y. Chevaleyre, and N. Maudet. The power of swap deals in distributed resource allocation. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, pages 625–633, 2015.
- [9] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.
- [10] D. K. Foley. Resource allocation and the public sector. *YALE ECON ESSAYS*, 7(1):45–98, 1967.
- [11] R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-2004)*, pages 125–131. ACM Press, 2004.
- [12] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, pages 873–880, 2010.
- [13] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14*, pages 675–692, New York, NY, USA, 2014. ACM.
- [14] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proceedings of the 1998 AAAI Spring Symposium on Satisficing Models*, 1998.
- [15] H. Steinhilber. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.