Accepted Manuscript

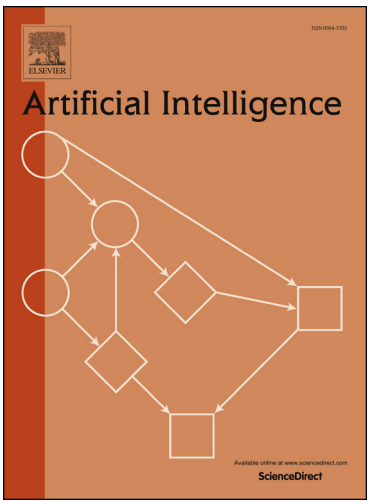

Distributed fair allocation of indivisible goods

Yann Chevaleyre, Ulle Endriss, Nicolas Maudet

PII: S0004-3702(16)30112-6
DOI: http://dx.doi.org/10.1016/j.artint.2016.09.005
Reference: ARTINT 2977

To appear in: *Artificial Intelligence*

Received date: 17 March 2016
Revised date: 23 September 2016
Accepted date: 26 September 2016

Please cite this article in press as: Y. Chevaleyre et al., Distributed fair allocation of indivisible goods, *Artif. Intell.* (2016), http://dx.doi.org/10.1016/j.artint.2016.09.005

Distributed Fair Allocation of Indivisible Goods[☆]

Yann Chevaleyre[a], Ulle Endriss[b], Nicolas Maudet[c]

[a] *LIPN, Université Paris-Nord, France*
[b] *ILLC, University of Amsterdam, The Netherlands*
[c] *Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris*

Abstract

Distributed mechanisms for allocating indivisible goods are mechanisms lacking central control, in which agents can locally agree on deals to exchange some of the goods in their possession. We study convergence properties for such distributed mechanisms when used as fair division procedures. Specifically, we identify sets of assumptions under which any sequence of deals meeting certain conditions will converge to a proportionally fair allocation and to an envy-free allocation, respectively. We also introduce an extension of the basic framework where agents are vertices of a graph representing a social network that constrains which agents can interact with which other agents, and we prove a similar convergence result for envy-freeness in this context. Finally, when not all assumptions guaranteeing envy-freeness are satisfied, we may want to minimise the degree of envy exhibited by an outcome. To this end, we introduce a generic framework for measuring the degree of envy in a society and establish the computational complexity of checking whether a given scenario allows for a deal that is beneficial to every agent involved and that will reduce overall envy.

Keywords: multiagent systems, multiagent resource allocation, fair division, negotiation, social networks

1. Introduction

The problem of fairly dividing a number of goods between several agents has been studied in a variety of settings. First, we may distinguish allocation problems for *divisible* and for *indivisible goods*. The literature on cake cutting, for instance, is concerned with divisible goods [3, 4, 5]. As for indivisible goods, we can distinguish *assignment problems* [6], where each agent can consume at most a single good, from more general settings where each agent may receive a set (or bundle) of goods [7, 8, 9]. If agents can receive sets of goods, and if their preferences over the goods are not additively separable, then fair division becomes a *combinatorial optimisation problem*. We may also distinguish whether or not to allow for *monetary side payments* to be added to the bundles allocated to the agents, and if so we have to decide what assumptions to make regarding the agents' appreciation of money (such as *quasi-linearity*, for instance). Finally, there are many different ways in which to interpret the term *fairness* itself [10]. For instance, we may be interested in equitable allocations, in proportional allocations, or in allocations where agents do not envy each other. In this paper, we adopt a model where sets of indivisible goods need to be allocated to a number of agents; agents express their preferences in terms of valuation functions over sets of goods; and side payments are possible and agents have quasi-linear preferences regarding money. We focus on two fairness criteria: first, we look for solutions that are *proportionally fair* (i.e., where the utility of each of the n agents is at least $1/n$th of the value she ascribes to the full set of goods) and, second, we look for *envy-free* solutions (where no agent envies any of the other agents). (These criteria will get formally defined in Sections 3 and 4.)

[☆]This paper extends work presented at the 20th International Joint Conference on Artificial Intelligence [1] and the 22nd AAAI Conference on Artificial Intelligence [2]. The first of these papers was co-authored with Sylvia Estivie and covers additional material not included here. We thank several anonymous reviewers for their helpful feedback. This work has been partly supported by COST Action IC1205 on Computational Social Choice as well as by project number ANR-14-CE24-0007-01 (CoCoRICo-CoDec) of the French National Research Agency.