Discrete Applied Mathematics 179 (2014) 54–68

Contents lists available at ScienceDirect

Discrete Applied Mathematics

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

Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods[☆]

Trung Thanh Nguyen [a], Jörg Rothe [b,*]

[a] Masdar Institute of Science and Technology, Abu Dhabi, United Arab Emirates

[b] Heinrich-Heine-Universität Düsseldorf, Germany

ARTICLE INFO

Article history:
Received 4 March 2014
Received in revised form 23 July 2014
Accepted 10 September 2014
Available online 5 October 2014

Keywords:
Resource allocation
Economically motivated agents
Approximability

ABSTRACT

Envy-freeness is a desirable criterion when one wishes to fairly distribute a finite set of goods among two or more agents. Unfortunately, allocations satisfying this criterion may not exist in the setting where the goods are assumed to be indivisible. In this case, it is useful to settle for allocations with envy as small as possible. Adapting the framework of Chevaleyre et al. (2007), we propose a multiplicative form of the degree of envy of a given allocation and then study the approximability of the corresponding envy minimization problems. We show that these problems are APX-hard to approximate in general, but admit an FPTAS for a fixed number of agents with additive utility functions. We also present a polynomial-time algorithm for the case when the number of agents is equal to the number of goods to be distributed. In addition, we study the problem of maximizing social welfare by the average Nash product. We provide a fast greedy approximation algorithm for this problem when the agents' utility functions are (sub)additive, and we design a PTAS for the case when all agents have the same additive utility function.

1. Introduction

The problem of fairly and efficiently allocating a finite set of goods to a number of agents has received growing interest in both economics and computer science in the last few decades, especially due to its wide range of potential applications, such as auctions, scheduling, network routing, airport traffic management, allocation of mineral riches in the ocean bed, the fair and efficient exploitation of earth observation satellites, greenhouse gas emissions reduction, and divorce settlement (see the survey of Chevaleyre et al. [10] for a detailed discussion of such applications). We study the allocation problem in the setting where goods are assumed to be indivisible and nonshareable, and agents to express their preferences over the goods by means of additive utility functions.

The goal is to find an allocation of goods to agents so as to satisfy certain desirable criteria, such as maximizing *egalitarian social welfare* (i.e., maximizing the utility of the agent that is worst-off) or *utilitarian social welfare* (i.e., maximizing the sum of all agents' individual utilities), or guaranteeing *envy-freeness* (i.e., no agent prefers another agent's bundle to her own),

☆ Preliminary versions of parts of this paper appear in the Proceedings of the *12th International Conference on Autonomous Agents and Multiagent Systems* Nguyen and Rothe (2013) [27], of the *6th International Workshop on Optimisation in Multi-Agent Systems* Nguyen and Rothe (2013) [26], and of the *3rd International Conference on Algorithmic Decision Theory* Nguyen and Rothe (2013) [28].

* Corresponding author. Tel.: +49 2118112188; fax: +49 2118111667.
E-mail addresses: ttnguyen@masdar.ac.ae (T.T. Nguyen), rothe@cs.uni-duesseldorf.de (J. Rothe).

http://dx.doi.org/10.1016/j.dam.2014.09.010

or satisfying *Pareto efficiency*.[1] In particular, the purpose of this paper is to investigate certain relaxed notions of envy-freeness and the notion of *(average) Nash social welfare* (see the work of Nash [21] and Kaneko and Nakamura [13]), which is measured as the (nth root of the) product of the agents' utilities, assuming there are n agents. In particular, we study the problem of computing an optimal allocation corresponding to these concepts from a computational point of view. To the best of our knowledge, only a few papers have attempted to tackle these problems in terms of computational complexity and approximation algorithms (see, e.g., [17,30,23,22,24]).

In Section 4, we consider envy-freeness: An allocation is said to be *envy-free* if no agent wants to swap her bundle of goods in that allocation with another agent. Unlike egalitarian social welfare, envy-freeness does not require interpersonal comparability of individual preferences. For the setting where goods are divisible, the envy-free allocation problem has been studied intensively during the last few decades under the name *cake cutting* (see, e.g., the textbooks by Brams and Taylor [7] and Robertson and Webb [31]). Although envy-free allocations are always possible when the goods are divisible, this is not the case for indivisible goods, assuming that all the goods must be assigned to the agents. For example, let us consider a quite simple scenario with only one good and many agents: Allocating the good to any agent will make the other ones envious. Therefore, it would make sense to pay attention to finding allocations whose envy is as small as possible.

There are various ways for defining the envy of an allocation. Chevaleyre et al. [11] proposed a framework for defining the "degree of envy" of an allocation based on the degree of envy among individual agents. Their approach can be seen as a generalization of that of Lipton et al. [17]. In more detail, assume that π_i and π_j are the bundles of goods assigned to, respectively, agent a_i and agent a_j in an allocation π. Agent a_i's envy regarding agent a_j's bundle is determined as $\max\{0, u_i(\pi_j) - u_i(\pi_i)\}$, where $u_i(\pi_i)$ and $u_i(\pi_j)$ are a_i's utility for the bundle π_i and π_j, respectively. Given the envy of agent a_i for each of the remaining agents, one can either sum these values up or take their maximum to obtain a_i's envy with respect to π. Finally, the envy of π can be computed by using the sum or the maximum operator again, over the envy of all agents with respect to π. Considering the optimization problems based on this measure of envy, a drawback of this approach is that, unless $P = NP$, there are no approximation algorithms for them, since the objective function might be zero (see the work of Lipton et al. [17]). We circumvent this by defining a similar notion of degree of envy in a multiplicative rather than additive way (see Section 4.1 for the formal definition).

Fairness and efficiency criteria cannot be met simultaneously in general; there is usually a trade-off between them (see, e.g., the book by Kaplow and Shavell [14]). For example, maximizing utilitarian social welfare in general does not give an optimal solution in terms of egalitarian social welfare, nor an envy-free allocation. Naturally, one wants to look for compromises between these criteria. The compromise we will focus on in Section 5 is maximizing social welfare by the (average) Nash product (which is sometimes also called "Bernoulli–Nash" social welfare).[2]

Like utilitarianism and egalitarianism, (average) Nash social welfare is a function of a society's individual utilities: It is defined as the (average of the) product of the individual agents' utilities. Why should this social welfare measure be studied? There are several reasons. First, the Nash product not only has a simple mathematical structure but it also takes into account both allocation efficiency and fairness. On the one hand, it has the (strict) monotonicity property of utilitarianism: Increasing any agent's utility leads to increasing the Nash product (provided that all agents have nonnegative utility,[3] which we do require for this measure of social welfare throughout this paper). On the other hand, the higher the Nash product, the more equal the individual agents' utilities are for an allocation—the Nash product may thus be seen as a measure of fairness. For example, if there are two agents and we have two alternative allocations whose utility vectors are $(2, 6)$ and $(4, 4)$, then utilitarian social welfare is indifferent between choosing any of these two allocations, whereas the Nash product will prefer the second one. Another important reason to study Nash social welfare is its desirable axiomatic properties, including so-called *independence of unconcerned agents*, the *Pigou–Dalton transfer principle*, and *independence of common utility scale* (see the work of Kaneko and Nakamura [13] and Moulin [19]). It is worth noting that while Nash social welfare is scale-independent, neither utilitarianism nor egalitarianism fulfills this axiomatic property. Further advantageous features of Nash social welfare can be found in the textbook by Moulin [20].

Having explained why both envy-freeness and Nash social welfare deserve to be carefully studied in the context of allocation of indivisible goods, let us finally mention that there might also be connections between these two criteria. For example, in the scenario when all agents have the same utility function, we suspect that each of the four envy minimization problems to be defined in Section 4.1 might be equivalent to the Nash social welfare maximization problem.

Organization. The rest of this paper is organized as follows. We first give a brief survey of previous related work and present our contributions in Section 2. In Section 3, we provide some background on the allocation problem for indivisible goods as well as the fundamentals of approximation algorithms needed to solve the problems defined in the following sections. In Section 4, we define the notion of degree of envy and present the (in-)approximability results for the corresponding envy minimization problems. Section 5 deals with the approximability of the problem of maximizing the average Nash product. Finally, Section 6 provides some conclusions and lists some open questions for future work.

1 An allocation is said to be *Pareto-efficient* (or *Pareto-optimal*) if no agent can improve her welfare without reducing another agent's welfare.

2 Another possibility for such a compromise would be to use the *leximin ordering* (see, e.g., Moulin [20]), which is a refinement of egalitarianism that, informally speaking, works by comparing first the utilities of the least satisfied agents, and when these coincide, comparing the utilities of the next least satisfied agents, and so on. While a leximin-optimal solution definitely maximizes the utility of the least happy agent, it is still Pareto-optimal.

3 Note that without this assumption the optimal solutions by the Nash product would not be continuous.