Envy-Ratio and Average-Nash Social Welfare Optimization in Multiagent Resource Allocation

(Extended Abstract)

Trung Thanh Nguyen and Jörg Rothe
Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany
{thanh, rothe}@cs.uni-duesseldorf.de

ABSTRACT

The resource allocation problem deals with distributing a number of indivisible, nonshareable resources among a set of agents so as to optimizing social welfare. Assuming all agents to have additive utility functions and focusing on two particular measures of social welfare, envy-ratio and average-Nash product, we investigate the two resulting optimization problems. We give the first hardness of approximation result for a factor better than $3/2$ for the problem of minimum envy-ratio, and we design an FPTAS for the case when the number of agents is fixed. For the special case when the number of agents and the number of resources are equal, we show that the problem is even solvable in polynomial time. Next, we propose the first approximation algorithm for maximizing the average-Nash product in the general case, and we prove that this problem admits a PTAS if all agents' utility functions are the same. Finally, we study the problem of how hard it is to design a truthful mechanism for these two optimization problems.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—Economics

General Terms

Economics, Theory

Keywords

Multiagent resource allocation, social welfare optimization, approximation algorithm, economically-motivated agents, auction and mechanism design

1. INTRODUCTION

A key task in multiagent resource allocation (MARA, for short) is to fairly distribute a given set of resources among a set of agents so as to optimizing social welfare. The survey by Chevaleyre et al. [3] provides many applications of this problem, ranging from computer science to economics and politics. In the standard model of MARA, we are given a set of agents and a list of indivisible, nonshareable resources. Every agent has individual preferences over the bundles of resources, which can be different from each other. The goal is to assign all resources to the agents (where each resource can be assigned to only one agent) so as to optimizing social welfare, which can be described by means of a *social welfare function* (SWF, for short). The most common SWFs, which have been studied intensely, include the *utilitarian SWF*, the *egalitarian SWF* (a.k.a. *Rawls's SWF*), and the *Nash (product) SWF*. Informally put, while utilitarian and Nash SWF are the (average) sum and the (average) product, respectively, of the agents' individual utilities, the egalitarian SWF gives the utility of an individual that is worst off in a given allocation. In this paper, we will focus on the (average) Nash SWF, which can be seen as a compromise between the (average) utilitarian and egalitarian SWFs. Indeed, with the utilitarian SWF it shares the monotonicity property, as in the mean increasing any agent's utility leads to an increase of the (average) Nash SWF. Moreover, the (average) Nash SWF provides a measure of fairness similarly to the egalitarian SWF, since the (average) Nash SWF increases when the differences between the agents' utilities decrease. This type of SWF is viewed as a useful measure of social welfare in economics, especially in ranking income distribution, which emphasizes the differences among the income gaps between the people in a country (see, e.g., [5, 2, 9]). Some other advantages of using this type of SWF are given by Moulin [6].

In some situations we wish to find an allocation that is *envy-free*, i.e., no agent evaluates another agent's bundle to be more valuable than her own. This is another useful criterion of fairness. Unfortunately, since all resources are indivisible and nonshareable and must be assigned completely to the agents, envy-free allocations do not always exist. For example, for an instance with two agents and only one resource that has positive utility for both agents, it is impossible to ensure envy-freeness: Whoever comes away empty-handed will be envious. A natural way to overcome this obstacle is to find allocations in which envy is as small as possible. Lipton et al. [4] suggested two minimization problems, one to minimize (total) envy and the other to minimize the envy-ratio. The second problem is more interesting and will be studied in this paper.

It is known that almost all social welfare optimization problems in MARA are NP-hard (see, e.g., [8, 4]). This motivates to study (in)approximability of these optimization problems.

2. PRELIMINARIES

Basics of Multiagent Resource Allocation

A MARA-setting is a triple $M = (A, R, U)$, where $A = \{a_1, \ldots, a_n\}$ is the set of *agents*, $R = \{r_1, \ldots, r_m\}$ is the set of indivisible and nonshareable *resources*, and every agent a_i is assumed to have an *additive utility function* $u_i \in U$ over the bundles of resources. Formally, each function u_i is a mapping from the power set of resources

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

2^R to $\mathbb{R}_+$, the set of nonnegative real numbers, such that for any bundle of resources $T \subseteq R$, $u_i(T) = \sum_{r_j \in T} u_i(r_j)$. We also assume that the empty bundle has always value zero for all agents. An *allocation* of resources to the agents is represented by a function $X : A \rightarrow 2^R$ such that $X(a_i)$ is the bundle assigned to agent a_i, and we require $X(a_i) \cap X(a_j) = \emptyset$ for all distinct $a_i, a_j \in A$, and $\bigcup_{a_i \in A} X(a_i) = R$. That is, each resource is assigned to at most one agent and all resources must be assigned completely to all agents.

DEFINITION 1. *For a MARA setting (A,R,U) and an allocation $X : A \rightarrow 2^R$, define the*

1. envy-ratio of X *as* $sw_{ev}(X) = \max_{1 \leq i,j \leq n} \left\{ 1, \frac{u_i(X(a_j))}{u_i(X(a_i))} \right\}$,
2. total-Nash (product) of X *as* $sw_N(X) = \prod_{i=1}^{n} u_i(X(a_i))$,
3. average-Nash (product) of X *as* $sw_{\overline{N}}(X) = \left(\prod_{i=1}^{n} u_i(X(a_i)) \right)^{1/n}$.

We model the two optimization problems as follows.

MIN-ENVY-RATIO	
Input:	A MARA setting $M = (A,R,U)$.
Task:	Find an allocation X that minimizes $sw_{ev}(X)$.

MAX-AVERAGE-NASH	
Input:	A MARA setting $M = (A,R,U)$.
Task:	Find an allocation X that maximizes $sw_{\overline{N}}(X)$.

Some Notions of Approximation Theory

DEFINITION 2. 1. *A* ρ-approximation algorithm *for an optimization problem is a polynomial-time algorithm that gives a feasible solution whose value is guaranteed to be within a factor of ρ of the optimum. The* approximation factor ρ *(a.k.a. the* performance guarantee*) of a ρ-approximation algorithm may depend on the input size; $\rho < 1$ for maximization problems and $\rho > 1$ for minimization problems.*

2. *An optimization problem L has a* polynomial-time approximation scheme *(*PTAS*) if for each ε, $0 < \varepsilon < 1$, there is a ρ-approximation algorithm for L, where $\rho = 1 - \varepsilon$ if L is a maximization problem, and $\rho = 1 + \varepsilon$ if L is a minimization problem. A* fully polynomial-time approximation scheme *(*FPTAS*) is a PTAS whose running time is bounded by a polynomial of the input size and of $1/\varepsilon$.*

3. RESULTS

By using a reduction from the well-known NP-complete problem EXACT-COVER-BY-3-SETS, we obtain the following inapproximability results:

THEOREM 3. MIN-ENVY-RATIO *is* NP*-hard to approximate within a factor better than $3/2$, even when the utility functions are restricted to the domain $\{0,1,3\}$.*

THEOREM 4. MAX-AVERAGE-NASH *cannot have a PTAS, unless* P = NP.

Turning now to approximability, whenever the number of agents is fixed, we can design an FPTAS for our two problems based on dynamic programming.

THEOREM 5. MIN-ENVY-RATIO *and* MAX-AVERAGE-NASH *admit an FPTAS for any fixed number of agents.*

In particular, if there are as many agents as resources, we can solving the problems efficiently by transferring them into a problem of finding a maximum matching on a weighted bipartite graph.

THEOREM 6. MIN-ENVY-RATIO *and* MAX-AVERAGE-NASH *can be solved in polynomial time when the number of agents and the number of resources are the same.*

Also by a matching technique, we provide a simple approximation algorithm for the problem of maximizing (total) average Nash social welfare.

LEMMA 7. MAX-AVERAGE-NASH *and* MAX-TOTAL-NASH *can be approximated to a factor of $1/(m-n+1)$ and of $1/(m-n+1)^n$, respectively.*

Theorem 8 is inspired by Lipton [4] and Alon et al. [1].

THEOREM 8. MAX-AVERAGE-NASH *has a PTAS for the case that all agents have the same utility for each of the resources.*

Regarding truthful mechanism design (see [7] for more details), we have the following results.

THEOREM 9. *There does not exist any truthful mechanism that computes an optimal allocation for* MAX-TOTAL-NASH *or for* MAX-AVERAGE-NASH.

THEOREM 10. *Let $0 < \varepsilon < 1$. There is no truthful mechanism that can yield an ε-approximation for* MAX-TOTAL-NASH *or for* MAX-AVERAGE-NASH.

THEOREM 11. *No truthful mechanism for* MIN-ENVY-RATIO *can have an approximation factor of $2 - \varepsilon$ for each fixed $\varepsilon > 0$.*

4. ACKNOWLEDGMENTS

This work was supported in part by DFG grant RO-1202/14-1, a DAAD grant for a PPP project in the PROCOPE programme, and a fellowship from the Vietnamese government. We thank the anonymous reviewers for their helpful comments

5. REFERENCES

[1] N. Alon, Y. Azar, G. Woeginger, and T. Yadid. Approximation schemes for scheduling. In *Proceedings of the 8th ACM-SIAM Symposium on Discrete Algorithms*, pages 493–500. Society for Industrial and Applied Mathematics, Jan. 1997.

[2] R. Basmann and D. Slottje. A new index of income inequality: The B measure. *Economics Letters*, 24(4):385–389, 1987.

[3] Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[4] 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*, pages 125–131. ACM Press, 2004.

[5] R. Moore. Ranking income distributions using the geometric mean and a related general measure. *Southern Economic Journal*, 63(1):69–75, 1996.

[6] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.

[7] N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the 31st ACM Symposium on Theory of Computing*, pages 129–140. ACM Press, May 1999.

[8] M. Roos and J. Rothe. Complexity of social welfare optimization in multiagent resource allocation. In *Proceedings of the 9th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 641–648. IFAAMAS, May 2010.

[9] D. Slottje, R. Basmann, and M. Nieswiadomy. On the empirical relationship between several well-known inequality measures. *Journal of Econometrics*, 42(1):49–66, 1989.