

Nash Social Welfare in Multiagent Resource Allocation

Sara Ramezani ^{ab}

Ulle Endriss ^a

^a *ILLC, University of Amsterdam, The Netherlands*

^b *CWI, Amsterdam, The Netherlands*

Abstract

We study different aspects of the multiagent resource allocation problem when the objective is to find an allocation that maximizes Nash social welfare, the product of the utilities of the individual agents. The Nash solution is an important welfare criterion that combines efficiency and fairness considerations. We show that the problem of finding an optimal outcome is NP-hard for a number of different languages for representing agent preferences; we establish new results regarding convergence to Nash-optimal outcomes in a distributed negotiation framework; and we design and test algorithms similar to those applied in combinatorial auctions for computing such an outcome directly.

This is an extended abstract of work presented in The Eleventh International Workshop on Agent Mediated Electronic Commerce [5].

1 Introduction

This work deals with multiagent resource allocation (MARA) with the aim of optimizing Nash social welfare. Multiagent resource allocation is an active field of research involving concepts and methods from various fields such as artificial intelligence, economics, and social choice theory. It has many applications, particularly with the rising use of software agents and e-commerce. The problem in MARA is that of dividing a limited amount of resources among a number of agents in a way that satisfies certain criteria.

We suppose that there is a finite number of agents and a finite number of indivisible and unsharable goods or resources. It is supposed that agents have preferences for various assignments of the resources to them, i.e. *allocations*, expressed by *utility functions*. Furthermore agents' utilities for each allocation depend only on the resources that they possess, and not on other agents' possessions.

The criteria to be satisfied usually involves efficiency. The most common efficiency criterion is *utilitarian* social welfare, the aim of which is to maximize the sum of the agents' utilities in the final allocation. This is equivalent to optimizing the auctioneer's revenue in the centralized version of the MARA problem, which is a combinatorial auction. Another approach can be to also consider fairness when assessing the collective welfare of the agents. There has been some work with this approach (e.g. [2]), but efficiency is usually considered as the sole collective welfare property to be satisfied.

Nash social welfare, on the other hand, considers maximizing the *product* of utilities of agents as desirable for the collective welfare of the agent society. Nash collective welfare is an important social welfare criterion studied extensively in the literature [3]. Its significance is that it takes both efficiency and fairness considerations into account, thus optimizing Nash social welfare can be a desirable criterion when fairness is an issue as well as efficiency or individual rationality of agents. It can also be characterized by some reasonable and desirable axioms which make it more appealing as a collective welfare function.

With this approach we address the MARA problem in two contexts: one in which agents can agree on an allocation by negotiating a series of deals, and another one where a central authority computes a Nash optimal allocation. We prove a number of theoretical results in the first and propose and experiment with a heuristic algorithm for solving the second. Details of proofs of all results can be found in [5] and [4].

2 Distributed Negotiation

In the first approach we consider a framework that has been studied previously [2, 6]. In this framework agents start with an initial allocation and proceed by negotiating *deals* between themselves. A deal is simply a pair of allocations, which means that agents go from the first allocation to the second. It has been shown that when such deals are individually rational, the resulting allocation after no more such deals are possible is optimal in terms of utilitarian social welfare [6]. Similar results for some other social welfare criteria are investigated in [2]. We consider the problem with the aim of optimizing Nash social welfare, which has not been studied before.

We define a *Nash deal* as a deal in which the product of the utilities of the agents involved in it (those whose bundle of assigned resources changes) increases. We show that when restricting agents to using Nash deals, which improve the local Nash welfare between them, the final allocation reached (when no such deal is further possible) will have maximal Nash social welfare globally. The Nash social welfare also increases in every step. However, the downside is that it may be necessary to use arbitrarily complex deals in such a negotiation sequence. Also, the deals could be complex enough to involve all agents in an inseparable fashion even when the utility functions of agents are restricted to the rather restrictive set of modular functions. This is not desirable because it may be hard for the agents to compute such deals in the negotiation process.

Furthermore, we show that while it is always possible to reach a Nash optimal allocation with at most one deal, if the agents are not able to compute a short path, there are cases in which the sequence of Nash deals may pass through every possible allocation, which is exponential in the number of resources.

3 Combinatorial Auctions

In the second approach we consider the centralized MARA problem which is equivalent to a *combinatorial auction* [1]. Here a central authority is responsible for computing the optimal allocation. Normally, the auctioneer in a combinatorial auction tries to maximize his revenue (supposing that the agents are willing to pay amounts proportional to their utilities for each bundle of goods) which is again equivalent to computing an allocation with maximal utilitarian social welfare. The problem of computing such an allocation is known as the *winner determination problem (WDP)*. We consider a variation in which the auctioneer tries to optimize Nash social welfare or solve the *Nash WDP*.

It is first shown that the Nash WDP is NP-complete when agents express their utilities in two different representation languages, namely the so-called XOR bidding language and a language based on weighted propositional formulas. Then an admissible heuristic (A^*) algorithm is designed for solving the Nash WDP when the aforementioned logic-based language is used. The algorithm has been implemented and some experiments have been run with it. It reduces the search space considerably in these experiments, but the resulting search tree still grows exponentially with a factor smaller than the complete tree. Hence the algorithm can be considered as a first step that shows that such approaches are possible for solving the Nash WDP, but more work would be needed to find efficient algorithms.

References

- [1] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.
- [2] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.
- [3] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- [4] S. Ramezani. Nash social welfare in multiagent resource allocation. Master's thesis, ILLC, University of Amsterdam, 2008.
- [5] S. Ramezani and U. Endriss. Nash social welfare in multiagent resource allocation. In *Proceedings of the 11th International Workshop on Agent-Mediated Electronic Commerce (AMEC-2009)*, May 2009.
- [6] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.