A fair payoff distribution for myopic rational agents

(Extended Abstract)

Stéphane Airiau
Institute for Logic, Language and Computation
University of Amsterdam
s.airiau@uva.nl

Sandip Sen
Computer Science department
University of Tulsa, Tulsa, USA
sandip@utulsa.edu

ABSTRACT

We consider the case of self-interested agents that are willing to form coalitions for increasing their individual rewards. We assume that each agent knows its individual payoff when a coalition structure (CS) is formed. We consider a CS to be stable if no individual agent has an incentive to change coalition from this CS.

When no stable CSs exist, rational agents will be changing coalitions forever. When stable CSs exist, they may not be unique, and choosing one over the other may give an unfair advantage to some agents. In addition, it may not be possible to reach a stable CS from any CS using a sequence of myopic rational actions. We propose a payoff distribution scheme that is based on the expected utility of a rational myopic agent (an agent that changes coalitions to maximize immediate reward) given a probability distribution over the space of coalition structures for selecting the initial CS. To compute this expected utility, we model the coalition formation process by a Markov chain. We recommend that agents share the utility from a social welfare maximizing CS proportionally to the expected utility they would receive if all agents acted in a myopic rational fashion. This scheme guarantees that agents receive at least as much as their expected utility from myopic behavior, which ensures sufficient incentives for the agents to use our proposed payoff distribution.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent system; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Social/Organizational Structure::Groups & Teams, Economics, theoretical, Markov chain, fairness, payoff distribution

1. INTRODUCTION

Cite as: A fair payoff distribution for myopic rational agents (Extended Abstract), Stéphane Airiau and Sandip Sen, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. XXX-XXX.

When agents jointly pick an alternative, a central problem is to select the "most desirable" alternative, as there may not be any alternative that satisfies all the agents at the same time. In this paper, an alternative is coalition structure (CS), i.e., a partition of the agents into coalitions. It is possible that for any CS, at least one agent has an incentive to change coalitions. Researchers have identified conditions for the existence of stable CS [1, 2, 3], but we want to provide a payoff distribution to promote stability even when no stable CS exists. The compromise we propose is based on the choice of an utilitarian social welfare maximizing CS and on allowing transfer of utility to provide a fair payoff to agents that induces stability.

We want the agent society to perform well as a whole. Hence our mechanism selects a social welfare maximizing CS. The computation of the side-payments to stabilize the CS is based on the expected utility of myopic and rational agents, i.e., agents that change coalition to maximize their immediate payoff). We believe that these expected utility values reflect the relative importance of agents in the coalition formation process as it is a global characteristic of the system, and hence payoff distributions based on these expected utility values can be fairn. Accordingly, we propose sharing the utility of a social welfare maximizing CS between agents proportionally to their expected utility values. Under the assumption that the initial CS is chosen at random, we show that each agent is better off following our protocol rather than acting in a myopic rational fashion.

2. PROBLEM DESCRIPTION

We consider a set N of n agents; N is also known as the $g$rand coalition. A coalition structure (CS) $s = \{\mathcal{S}_1, \cdots, \mathcal{S}_m\}$ is a partition of N, where $\mathcal{S}_i$ is the i^{th} coalition of agents with $\cup_{i\in[1..m]}\mathcal{S}_i = N$ and $i \neq j \Rightarrow \mathcal{S}_i \cap \mathcal{S}_j = \emptyset$. $\mathscr{S}$ is the set of all CSs. The coalition of agent i in s is noted as $s(i)$. We consider that an agent i has a valuation funtion $v_i(s)$. Note that the valuation function not only depends on the other agents present in the coalition but also on the other coalitions present in the population. We further assume that there is no coordinated change of coalitions; only one agent at a time can change coalition. These assumptions prevent uncertainties about the state of the CS. After a change of CS, it is possible, if not likely, that another agent changes its coalition, leading to a different CS. As it is computationally expensive to perform multi-step look ahead because of the large state space, we consider myopic agents that change coalitions to maximize their immediate reward (agents act sincerely but not strategically). We believe it is reasonable to assume that current members can control when other agents can join a coalition. Hence we require that all members of a coalition must agree to accept a new member and, if some member i refuses, i can veto the transition.

Our main stability criterion is individual stability [3]: a CS s is *individually stable* when there is no agent i that can change coalition from $s(i)$ to another existing coalition $C \in s$ so that the payoff of agent i increases and the payoff of no member of C decreases. We also require that members of a coalition cannot prevent a member to leave, even if some of the remaining members lose utility.

3. A FAIR PAYOFF DISTRIBUTION FOR MYOPIC RATIONAL AGENTS

For the society to operate efficiently, we want the agents to form a stable CS. What we propose to the agents is to sign a binding contract to form a particular CS $s^\star$ and to receive a particular payoff u. If one agent does not want to sign the contract, the agents must find a stable CS on their own, for example using a random walk in the space of CS. We propose that the agents form a CS $s^\star$ that maximizes social welfare. However, $s^\star$ may not be stable as some agents may have incentives to change coalition from $s^\star$, hence we want to share the valuation $v(s^\star)$ so as to provide some fairness to the agents. On one hand, we want the entire population of agents to cooperate and work together, which suggests the use of the grand coalition. On the other hand, we want to use the synergy between the agents, and thus form a CS that maximizes social welfare.

The utility function, over the whole state space, reflects the relative importance of an agent in the society. A first candidate for sharing utility is to share $v(s^\star)$ proportionally to the average utility over all the CSs. This assumes that each CS is equally important and is unlikely to be true in general. If we assume any CS is likely to be the initial CS, the coalition formation process, driven by myopic rational agents, can be modeled by a Markov chain and, by studying the long term behavior of the chain, we can compute the expected utility of the agents under our hypothesis, i.e., when agents are myopic rational and when members of a coalition can veto the entrance of new members. The expected utility is a neutral and fair metric to determine and compare the strength of each agent in the coalition formation process.

The computation of the expected values for the agents assumes that the initial CS is chosen uniformly over the space of CSs, in other word, it is not biased by the initial CS[1]. We consider that this value represents the strength of an agent given the valuation function. An agent with high expected utility should obtain a larger payoff than agents with lower expected utility. We propose a distribution of $v(s^\star)$ to all agents proportional to the expected payoff of the agents, i.e., we prescribe the payoff to agent i to be

$$u_i = \frac{E(v_i)}{\sum_{j \in N} E(v_j)} v(s^\star),$$

where $E(v_i)$ is the expected value of agent i.

First, note that the payoff distribution is not according to the actual coalitions present in $s^\star$. In other words, it is possible that for each coalition $\mathcal{C} \in s^\star$, $\sum_{i \in \mathcal{C}} u_i \neq \sum_{i \in \mathcal{C}} v_i(s^\star)$. Then, note that u_i is guaranteed to be at least as good as $E(v_i)$, as shown by Property 1 (the proof is omitted for lack of space). In addition, the payoff distribution is Pareto Optimal as we share the value of a social welfare maximizing CS (if an agent gets more utility, at least another agent must lose some). We believe that these incentives are sufficient for the agents to accept our proposed value.

PROPERTY 1. $u_i = \frac{E(v_i)}{\sum_{j \in N} E(v_j)} v(s^\star) \geq E(v_i)$, *i.e., the payoff of an agent is at least as good as the utility that a myopic ration agent would get on average.*

Another important question is whether the payoff distribution u_i is individually rational: is an agent guaranteed to get as much as when an agent is forming a singleton coalition? The minimum payoff an agent can guarantee for itself is $r_i = \min_{s \in \mathscr{S},\, \{i\} \in s} v_i(s)$ [4]. As shown in Proposition 2, the answer is positive, but only when the minimum payoff occurs when an agent forms a singleton coalition. This is the most likely case in practice: the worst case scenario should be when it forms a singleton coalition and when all other agents in the population try to minimize its payoff.

PROPERTY 2. *If* $(\forall s \in \mathscr{S})\ v_i(s) \geq r_i = \min_{s \in \mathscr{S} | \{i\} \in s}$, *then* $u_i \geq r_i$, *i.e., the payoff distribution* u_i *is individually rational.*

4. CONCLUSION

Myopic rational agents who receive a private payoff that depends on the CS may never reach an agreement on the CS to be formed. We designed a protocol that computes a payoff distribution where agents 1) form the optimal CS, which makes the multiagent system efficient from the viewpoint of a system designer and 2) are guaranteed to have at least the expected utility from a process where each agent would change coalition to maximize its immediate reward. The protocol assumes that 1) the valuation function provides a payoff for each individual agent given a CS and 2) the agents are myopically rational. The payoff function we propose is based on the value of a social welfare maximizing CS and on the expected utility of the agents if they try to change coalitions to maximize their immediate reward. We argue that this is a fair distribution as the payoff obtained by an agent reflects the behavior of the agents over the entire space of CSs, i.e., it is a global property of the valuation function.

In this paper, we made the myopic assumption as it appears computationally expensive to perform multi-steps look-ahead. We would like to study, in the case of complete information if agents can consider strategic actions. We would like to study the existence of a farsighted equilibrium, which could strengthen the notion of fairness.

5. REFERENCES

[1] J. Alcalde and A. Romero-Medina. Coalition formation and stability. *Soc Choice Welfare*, 27(2):365–375, 2006.

[2] S. Banerjee, H. Konishi, and T. Sönmez. Core in a simple coalition formation game. *Soc Choice Welfare*, 18(1):135–153, 2001.

[3] A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Game Econ Behav*, 38(2):201–230, 2002.

[4] D. Ray and R. Vohra. A theory of endogenous coalition structures. *Game Econ Behav*, 26(2):286–336, 1999.

[1]If the coalition formation process has a bias, it is easy to incorporate it in the computations.