ECAI 2016
G.A. Kaminka et al. (Eds.)
© 2016 The Authors and IOS Press.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/978-1-61499-672-9-1606

Increasing Coalition Stability in Large-Scale Coalition Formation with Self-Interested Agents

Pavel Janovsky[1] and **Scott A. DeLoach**[1]

Abstract. In coalition formation with self-interested agents both social welfare of the multi-agent system and stability of individual coalitions must be taken into account. However, in large-scale systems with thousands of agents, finding an optimal solution with respect to both metrics is infeasible.

In this paper we propose an approach for finding coalition structures with suboptimal social welfare and coalition stability in large-scale multi-agent systems. Our approach uses multi-agent simulation to model a dynamic coalition formation process. Agents are allowed to deviate from unstable coalitions, thus increasing the coalition stability. Furthermore we present an approach for estimating coalition stability, which alleviates exponential complexity of coalition stability computation. This approach is used for estimating stability of multiple coalition structures generated by the multi-agent simulation, which enables us to select a solution with high values of both social welfare and coalition stability. We experimentally show that our algorithms cause a major increase in coalition stability compared to a baseline social welfare-maximizing algorithm, while maintaining a very small decrease in social welfare.

1 Introduction

Coalition formation is a process of grouping of agents into *coalitions* in order to increase the agents' cooperation. A goal of coalition formation is often to increase social welfare of the multi-agent system, which can generate unrealistic solutions if the agents prefer their own profit to the global social welfare. These self-interested agents would deviate from the computed social welfare-maximizing coalitions.

In coalition formation with self-interested agents, *stability* of the coalitions, which measures the coalition's ability to de-incentivize any sub-coalition of agents from leaving the coalition, must be addressed as a concept that along with the social welfare influences the coalition formation algorithms and solutions. Coalition stability is addressed in game theoretical literature mainly through the concept of a *core*, which is a set of allocations to the agents in a coalition, such that these allocations cannot be improved upon by allocations to a subset of these agents. While the *core* is a strong concept, its computation in a setting where coalition values are generated by general polynomial-time functions requires an evaluation of all $2^{|C|}$ possible sub-coalitions of each coalition C containing $|C|$ agents. In this setting even determining whether the *core* is non-empty is $\Delta_2^P - complete$ [3]. This complexity makes the use of the *core* in large-scale systems infeasible. Therefore instead of the *core* we approach coalition stability using multi-agent simulation. Instead of looking for stable distribution of the coalition value to the agents, we specify an allocation scheme beforehand and let the agents utilize this information to choose more stable coalitions.

Specifically, the contributions of this paper are the following:

1. *An algorithm for large-scale coalition formation that increases coalition stability by allowing agents to deviate.*
2. *An approach for selecting sub-optimal solutions based on their social welfare and coalition stability.*

2 Problem Statement

We study the coalition formation problem, in which agents $a_1, a_2, \ldots, a_n \in A$ form coalitions C_i such that each agent belongs to exactly one coalition. We assume that the agents have full information about each others' states. A coalition structure CS is a set of all coalitions C_i that the agents formed. The task is to find a coalition structure that maximizes its social welfare as well as its stability.

We represent the social welfare by a gain metric defined in [4] as $g(CS) = \frac{1}{n} \cdot (v(CS) - v(CS_0))$, where $v(CS)$ is a value of a coalition structure CS, which is a sum of coalition values $v(C)$ assigned by a polynomial-time function, and CS_0 denotes the coalition structure of singleton coalitions. The gain shows an average benefit of an agent in coalition formation.

Self-interested agents maximize their own profit, which we define as marginal contribution of an agent to a coalition at the time of entry. [1] describes games that use this profit as Labor Union games.

Since finding coalition stability is computationally expensive, we approximate it by $stability_\alpha$. We say that a coalition C is $\alpha-$stable if no sub-coalition D with $\langle 1, \alpha \rangle$ members can be formed in which some agents would benefit more and no agent would benefit less than in C. We define $stability_\alpha$ of a coalition structure in terms of α as

$$stability_\alpha(CS) = \frac{|\alpha\text{-stable coalitions in CS}|}{|CS|}, \quad (1)$$

where $|CS|$ denotes the number of coalitions in CS. With increasing α $stability_\alpha$ approaches the true stability of CS, which we define as the ratio of stable coalitions in CS. Since $stability_\alpha$ is non-increasing w.r.t. α, it is an upper estimate of this true stability of CS.

Finally, we use the price of stability $PoS(CS_{sw}, CS_{sa}) = g(CS_{sw})/g(CS_{sa})$ to show the ratio between the gain of social welfare maximizing solutions CS_{sw} and the gain of solutions reached by behavior of self-interested agents CS_{sa}.

3 Methodology

We find solutions to coalition formation using multi-agent simulation. We extend a multi-agent simulation framework for large-scale coalition formation proposed in [4], in which the agents maximize

[1] Department of Computer Science, Kansas State University, USA, email: {janovsky, sdeloach}@ksu.edu

the social welfare. In that framework the agents use strategies to decide about leaving their coalitions and joining new coalitions. The coalitions are evaluated by a polynomial-time valuation function. This process repeats in an iterative fashion, resulting in an agent-driven search of the coalition structures state-space. While [4] shows almost-optimal performance in small-scale scenarios and stable gain in large-scale scenarios, it does not consider stability of the solutions.

In order to increase stability of coalition structures we extend the algorithm from [4] by allowing the agents to create more stable sub-coalitions within their coalition by the process of deviation, and by selecting the best solution out of the pool of solutions generated by the simulation with respect to both social welfare and stability.

3.1 Deviation

Deviation guides the search towards more stable coalition structures by allowing agents to leave their current coalition along with other agents from the same coalition. A sub-coalition of agents can deviate from its coalition if no agent loses profit by deviation and at least one agent gains profit. Considering all $2^{|C|-1}$ possible sub-coalitions that an agent can be part of is infeasible, therefore the agents use a heuristic to guide their search. Some possible heuristics are adding agents to the sub-coalition in order of increasing and decreasing profit, and in random order. Our experiments showed that most stable coalitions were found using the increasing profit heuristic.

Deviation is performed in our model after the agents decide on leaving and joining coalitions. Each iteration of the simulation therefore consists of two steps: social welfare maximization by leaving and joining coalitions, and stability maximization by deviation.

3.2 Solution selection

An advantage of using multi-agent simulation for coalition formation is the fact that it creates a pool of solutions encountered during the search. At the end of the simulation, [4] selects from this pool a solution that maximizes the gain. We propose to select a solution based on both gain and stability metrics. However, computing stability of a coalition structure is computationally expensive, therefore we use $stability_\alpha$ to estimate the true stability of the solutions.

We compute $stability_\alpha$ in an iterative fashion for increasing $\alpha \in \langle 1, \alpha_{max} \rangle$. We only have to determine whether a coalition is α-stable if it is $(\alpha-1)$-stable. We mark a coalition C α-stable if in all permutations of all combinations of α agents from C some agents lose or no agent gains profit[2]. We then calculate $stability_\alpha$ using Eq. 1.

After $stability_\alpha$ of all coalition structures is computed, a multi-criteria optimization is used to select a best coalition structure based on its gain and $stability_\alpha$. Common approaches of multi-criteria optimization are finding Pareto optimal solutions and designing a fitness function. We use a simple fitness function that assigns a same weight to both social welfare and coalition stability.

4 Experimental Analysis

We tested our algorithm in a collective energy purchasing scenario from [6], which models agents as households that buy electricity based on their requested daily energy profiles. Electricity can be bought at spot markets based on current demand, and at cheaper forward markets based on demand prediction. Agents form coalitions in order to make their aggregate energy profiles more predictable so they could exploit the reduced prices of the forward market. We used a dataset of daily energy profiles of households in Portugal [5]. To prevent the trivial grand coalition from being the optimal solution, we use a coalition size penalty $\kappa = min(-|C|+\mu, 0)^\gamma$ with $\mu = 10$ and $\gamma = 1.1$. Agents move to new coalitions using *local search strategy* [4], which is a best response strategy with a random element. Results are averaged over 10 random runs of our algorithm.

We compared results of our algorithms with the baseline algorithm [4] using the $stability_\alpha$ and *price of stability* metrics, as shown in Table 1, in which the first line represents the baseline algorithm. Table 1 shows that our algorithms increase significantly the coalition stability, while introducing a necessary, but very low, price of stability. Combination of deviation and solution selection yields the highest coalition stability.

Table 1: Trade-off between average stability and average price of stability achieved by our algorithms with $\alpha = 4$.

Algorithm		Results	
Deviation	**Solution selection**	**Average** $stability_\alpha$	**Average** PoS
NO	NO	0.5538	-
YES	NO	0.7857	1.0193
NO	YES	0.7121	1.0029
YES	YES	0.8363	1.0347

5 Conclusion

Algorithms that find stable coalition structures are often proposed for problems that restrict the valuation functions or scale, such as [1, 2]. Practical aspects of the high complexity of finding stable coalitions for large-scale multi-agent systems are often not considered.

In this work we proposed an approach for increasing coalition stability in large-scale coalition formation with self-interested agents and arbitrary valuation functions. We modeled agent behavior using multi-agent simulation, in which we allowed the agents to choose profitable coalitions and deviate from unstable coalitions. At the end of the simulation, we selected a solution out of a pool of generated coalition structures based on its social welfare and stability. We experimentally showed that our approach is able to increase the stability of the solutions in a real-world scenario. We also showed that the necessary price for this increase in stability that our algorithm incurs to the social welfare is very low.

REFERENCES

[1] John Augustine, Ning Chen, Edith Elkind, Angelo Fanelli, Nick Gravin, and Dmitry Shiryaev, 'Dynamics of profit-sharing games', *IJCAI'11: 21st International Joint Conference on Artificial Intelligence*, (2011).

[2] Filippo Bistaffa and Alessandro Farinelli, 'A fast approach to form core-stable coalitions based on a dynamic model', *2013 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2013*, (2013).

[3] Gianluigi Greco, Enrico Malizia, Luigi Palopoli, and Francesco Scarcello, 'On the complexity of the core over coalition structures', *Twenty-Second International Joint Conference on Artificial Intelligence*, (2011).

[4] Pavel Janovsky and Scott A. DeLoach, 'Multi-agent simulation framework for large-scale coalition formation', *2016 IEEE/WIC/ACM International Conference on Web Intelligence*, (2016).

[5] M. Lichman. UCI machine learning repository, 2013. Available at https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014.

[6] Meritxell Vinyals, Filippo Bistaffa, Alessandro Farinelli, and Alex Rogers, 'Coalitional energy purchasing in the smart grid', *2012 IEEE International Energy Conference and Exhibition, ENERGYCON 2012*, (2012).

[2] All permutations must be considered because the order in which agents join coalitions determines their profit