Arbitration and Stability in Cooperative Games in Overlapping Coalitions

Yair Zick
School of Physical and Mathematical Sciences
Nanyang Technological University, Singapore
mail: yair0001@ntu.edu.sg

1 Introduction

Consider the following scenario: a group of agents is tasked with completing some projects; the agents divide into groups, and using the resources available to each group, agents generate profits, which must in turn be divided among group members. Cooperative game theory [Peleg and Sudhölter, 2007] studies such scenarios; formally, Given a set of agents $N = \{1, \dots, n\}$, the value of each subset S of N is given by a function $v : 2^N \to \mathbb{R}$. Agents first form a *coalition structure* CS by partitioning into disjoint sets; then, the value of each subset $S \in CS$ is divided among the members of S. Such payoff divisions are also called *imputations*. Given a game $\mathcal{G} = \langle N, v\rangle$, a *solution concept* for $\mathcal{G}$ is a set of imputations that share some desirable properties; for example, the *core* of a game $\mathcal{G}$ is the set of all payoff divisions such that for all $S \subseteq N$, the total payoff to S is at least $v(S)$. That is, the core is the set of all *stable* payoff divisions, from which no subset of agents would want to deviate.

Classic cooperative game theory assumes that when agents form coalition structures, each agent is a member of only one coalition. In Overlapping Coalition Formation (OCF) games [Chalkiadakis *et al.*, 2010], agents can participate in several coalitions. Each agent $i \in N$ controls some finite resource such as time, computational power, or money. The key feature of OCF games is that unlike classic cooperative games, agents are allowed to concurrently commit resources to several coalitions. Thus, a coalition is no longer a subset of N, but rather a *vector* $\mathbf{c}$ in $[0, 1]^n$, where the i-th coordinate of $\mathbf{c}$, c^i, denotes how much of i's resource is devoted to $\mathbf{c}$. The valuation function v is now from $[0, 1]^n$ to $\mathbb{R}$, rather than from 2^N to $\mathbb{R}$. Under this setting, a coalition structure CS is a list of vectors in $[0, 1]^n$, $(\mathbf{c}_1, \dots, \mathbf{c}_k)$, and its value is simply $\sum_{j=1}^k v(\mathbf{c}_j)$. Having formed CS, agents must divide the payoffs from CS in some manner; such a payoff division, $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_k)$, consists of vectors $\mathbf{x}_j$, such that $\sum_{i=1}^n x_j^i = v(\mathbf{c}_j)$. Similarly to the non-overlapping setting, if $c_j^i = 0$, i.e. agent i does not contribute to $\mathbf{c}_j$, then i may not receive any payoff from $\mathbf{c}_j$. Those agents for which $c_j^i > 0$ are called the *support* of $\mathbf{c}_j$. The pair $(CS, \mathbf{x})$ is called an *outcome* of $\mathcal{G}$.

As noted in [Chalkiadakis *et al.*, 2010], stability in OCF games is a complicated matter; when deviating from $(CS, \mathbf{x})$, a set S may abandon some, but not all of the coalitions it is involved in. The main issue is whether S gets to keep its payoffs under $(CS, \mathbf{x})$ from coalitions that are unaffected by the deviation. [Chalkiadakis *et al.*, 2010] introduce three possible reactions to deviation: the *conservative*, *refined*, and *optimistic*. Under the conservative reaction, S may expect no payoffs from any coalition; like in the non-overlapping case, it assumes that it is “on its own” if it deviates; under the refined reaction, S may expect payoff from all coalitions that were not changed by the deviation; finally, under the optimistic, S may still receive payoff from a coalition $\mathbf{c}_j$, if it can reduce its contribution to $\mathbf{c}_j$ while still paying all agents in $N \setminus S$ the same amount they got from $\mathbf{c}_j$ under $(CS, \mathbf{x})$.

2 Our Contribution: Arbitration Functions

We propose a general model for the study of stability in the OCF setting (see [Zick and Elkind, 2011]. Reaction to deviation is described by an *arbitration function* $\mathcal{A}$, whose input is an outcome $(CS, \mathbf{x})$, a deviating set S, and S's deviation from $(CS, \mathbf{x})$; $\mathcal{A}$'s output is a value $\alpha(\mathbf{c})$ specifying how much is the coalition $\mathbf{c} \in CS$ willing to give S, given its deviation. $\alpha(\mathbf{c})$ does not have to depend only on the effect S had on $\mathbf{c}$; it is possible that members of $\mathbf{c}$ are aware of global effects to the outcome. For example, $\alpha(\mathbf{c})$ is 0 if some agent in the support of $\mathbf{c}$ was hurt by S in some other coalition $\mathbf{c}' \in CS$.

Using this extension of the OCF model, we fully characterize arbitrated solution concepts and their properties. Our characterizations hold under minimal assumptions on the valuation function v and the arbitration function $\mathcal{A}$. We describe the arbitrated core, as well as other solution concepts for OCF games such as the nucleolus, bargaining set, and Shapley value; we show that the solution concepts we define share many of the properties of their non-OCF counterparts. For example, the arbitrated nucleolus is never empty and that it is always in the core if the latter is not empty (for a fixed arbitration function). We also show that the OCF Shapley value can be derived using an axiomatic approach. However, different axiomatic assumptions lead to two different values, which are the unique values which satisfy these axioms.

2.1 Computational Aspects of OCF games

There is a well established body of literature studying computational aspects of cooperative games (for a detailed literature review see [Chalkiadakis *et al.*, 2011]). [Chalkiadakis *et al.*, 2010] study some computational issues in OCF games, but they limit their attention to a class of OCF games called *threshold task games*. We study computational aspects of games with overlapping coalitions in [Zick *et al.*, 2012a], with the additional assumption that agent resources are given by integer weights, and that they may only allocate an integer amount of resources to a coalition. This is a natural discretization of the classic OCF model. Given a discrete OCF game $\mathcal{G} = \langle N, v \rangle$ and an arbitration function $\mathcal{A}$, we analyze the computational complexity of optimization and stability in OCF games, e.g. computing an optimal coalition structure, deciding the most a set can get by deviating, and deciding whether a given outcome is stable.

Perhaps unsurprisingly, such questions turn out to be NP-hard. We then turn to exploring what underlying structural properties of OCF games induce NP-hardness. It turns out that intractability stems to some extent from agents having large weights, but to a greater degree it stems from complex agent interaction. We show that if one assumes that agents have polynomially bounded weights and interactions are simple, then all above questions can be answered in polynomial time. Interactions must be simple in two respects. First, agents must not be allowed to form coalitions with whoever they wish; we show that if agents form a hierarchical interaction structure (i.e. a tree), or have an interaction structure that is nearly hierarchical (i.e. has a bounded treewidth), then all above questions can be answered in polynomial time. Second, agents' reaction to deviation must be local in nature. Recall that given an outcome $(CS, \mathbf{x})$, a set S and S's proposed deviation, the arbitration function $\mathcal{A}$ needs to assign a value $\alpha(\mathbf{c})$ for each coalition $\mathbf{c}$ in CS. $\alpha(\mathbf{c})$ can depend, in general, on S's effect on coalitions other than $\mathbf{c}$. However, we show that if $\mathcal{A}$ allows this (a behavior type which we term *global*), then one cannot compute the most S can get by deviating in polynomial time, unless P equals NP. Thus, it is important to assume that decisions on how much should S get from $\mathbf{c}$ should depend solely an S's effect on $\mathbf{c}$.

2.2 Characterizing Stable OCF Games

The arbitrated core of an OCF game is an appealing solution concept; however, it is often the case that it is empty. The objective of [Zick *et al.*, 2012b] is to provide characterizations of stable games, and offer sufficient conditions for arbitrated core non-emptiness. Given an OCF game $\mathcal{G} = \langle N, v \rangle$, we can construct an non-OCF game $\widehat{\mathcal{G}} = (N, U_v)$ where U_v is a function on subsets of N, with $U_v(S)$ equaling the most that S can make on its own. In a sense, $\widehat{\mathcal{G}}$ can be seen as a discrete, optimized version of $\mathcal{G}$. We show that $\mathcal{G}$ has a non-empty conservative core if and only if $\widehat{\mathcal{G}}$ has a non-empty core. Moreover, we show that our condition for conservative core non-emptiness implies the convexity condition stated in [Chalkiadakis *et al.*, 2010]. We also characterize coalition structures that can be stabilized w.r.t. the refined arbitration function; our result is similar to the Bondareva-Shapley [Shapley, 1967] characterization of stable non-OCF games. Using this characterization, we identify a sufficient condition on v which ensures that the refined core is not empty. Finally, we introduce a class of games which are guaranteed to have a non-empty optimistic core, which we call Linear Bottleneck Games (LBGs). Briefly, an LBG is described by a list of tasks $\mathcal{T} = (T_1, \dots, T_k)$, where each task requires the participation of some set of agents $A_j \subseteq N$; if the least contribution of the agents in A_j is x, then the profit generated by T_j is xp_j, where p_j is the value of T_j. These games model a variety of optimization problems, such as multicommodity flow games. Using linear programming techniques, we show that LBGs have a non-empty optimistic core.

3 Discussion

We introduce a general framework for handling deviation in OCF games, and analyze algorithmic and game-theoretic properties of the resulting model. Our work can be extended in several interesting ways. First, while we provide exact algorithms for computing solution concepts in OCF games, it would be useful to study approximation algorithms for finding coalition structures and stable payoff divisions in OCF games. Second, while we identify properties of OCF games that ensure stability for a given arbitration function, one can alternatively fix an OCF game $\mathcal{G}$, and identify arbitration functions that ensure that $\mathcal{G}$ is stable. Finally, while our definition of arbitration functions is applied only to a specific type of strategic agent interaction, namely cooperative games with overlapping coalitions, we believe that our approach is much more general. The effect that non-deviators have on the desirability of deviation can be studied in any agent setting where agents may still interact with non-deviators after their deviation. This occurs, for example, in congestion games, matching markets, and collaboration networks. It would be interesting to see what results would our methodology yield in such settings.

References

[Chalkiadakis *et al.*, 2010] G. Chalkiadakis, E. Elkind, E. Markakis, M. Polukarov, and N.R. Jennings. Cooperative games with overlapping coalitions. *JAIR*, 39:179–216, 2010.

[Chalkiadakis *et al.*, 2011] G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan and Claypool, 2011.

[Peleg and Sudhölter, 2007] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*. Springer, Berlin, second edition, 2007.

[Shapley, 1967] L.S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14(4):453–460, 1967.

[Zick and Elkind, 2011] Y. Zick and E. Elkind. Arbitrators in overlapping coalition formation games. In *AAMAS'11*, pages 55–62, 2011.

[Zick *et al.*, 2012a] Y. Zick, G. Chalkiadakis, and E. Elkind. Overlapping coalition formation games: Charting the tractability frontier. In *AAMAS'12*, pages 787–794, 2012.

[Zick *et al.*, 2012b] Y. Zick, E. Markakis, and E. Elkind. Stability via convexity and lp duality in ocf games. In *AAAI'12*, pages 1506–1512, 2012.