Stability and Arbitration in Cooperative Games with Overlapping Coalitions

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

1. INTRODUCTION

We begin by describing two important theoretical notions related to our body of work.

1.1 Classic Cooperative TU Games

In classic cooperative games with *transferable utility* [4] (TU games), there is a set of agents $N = \{1, \dots, n\}$, where each subset S of N has some value $v(S)$. The goal of the agents is to first form a *coalition structure* by partitioning N into disjoint sets; second, the value of each set in the partition is divided among its members. The payoff division $\mathbf{x} = (x_1, \dots, x_n)$, often referred to as an *imputation*, is then analyzed. It is often desirable that $\mathbf{x}$ maintains some notion of stability; over the years, several methods of payoff distribution, or *solution concepts*, have been suggested. Formally, given a game $\mathcal{G} = (N, v)$, a solution concept $SC(\mathcal{G})$ is a set of imputations; if $SC(\mathcal{G})$ is a singleton for all games, SC is called a *value*. One of the most popular solution concepts is the *core*. We say that an imputation $\mathbf{x}$ is in $Core(\mathcal{G})$ if for all sets $S \subseteq N$, $\sum_{i \in S} x_i \geq v(S)$. Other popular solution concepts are the Shapley value, the bargaining set, and the nucleolus.

1.2 OCF Games

Overlapping Coalition Formation (OCF) games [2] are cooperative games where 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, where an agent must devote all of his resources to a coalitions, agents are allowed to contribute only some of their effort to a coalition. 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 $\mathbb{R}$. 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 $\mathbf{c}_1, \dots, \mathbf{c}_k$ 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 [2], 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. In [2] the authors 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; 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. ARBITRATED OCF GAMES

In the paper [7], we propose a general model for the study of stability in the OCF setting. 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 ρ_j specifying how much is the coalition $\mathbf{c}_j$ willing to give S, given its deviation. ρ_j does not have to depend only on the effect S had on $\mathbf{c}_j$; it is possible that members of $\mathbf{c}_j$ are aware of global effects to the outcome. For example, ρ_j is 0 if some agent in the support of $\mathbf{c}_j$ was hurt by S in some other coalition $\mathbf{c}_{j'}$.

Using this extension of the OCF model, we proceed to 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, 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.

3. 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 [3]). [2] 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 [6]. Given a game $\mathcal{G} = (N, v)$ and an arbitration function $\mathcal{A}$,

1. OPTVAL: Can we find an optimal coalition structure in polynomial time?

2. MAXDEV: Given an outcome $(CS, \mathbf{x})$, can we compute the most a set $S \subseteq N$ can get by deviating?

3. INCORE: Given an outcome $(CS, \mathbf{x})$, is $(CS, \mathbf{x})$ in the core under $\mathcal{A}$?

4. ISSTABLE: Given a coalition structure CS, can we find $\mathbf{x}$ such that $(CS, \mathbf{x})$ is in the core under $\mathcal{A}$?

Our goal is to analyze the computational complexity of the above questions. Unfortunately, even OPTVAL is NP-hard. We then turn to exploring what structural issues in 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 ρ_j for each coalition $\mathbf{c}_j$ in CS. ρ_j can depend, in general, on S's effect on coalitions other than $\mathbf{c}_j$. 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 $\mathcal{A}$ has *local* behavior, i.e. its decision on how much should S get from $\mathbf{c}_j$ should depend solely an S's effect on $\mathbf{c}_j$.

4. 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 [8] is to provide characterizations of stable games, and offer sufficient conditions for arbitrated core non-emptiness. Indeed, [2] provide an initial characterization of outcomes that are stable w.r.t. the conservative arbitration function via a notion of *convexity* for the characteristic function v. Given an OCF game $\mathcal{G} = (N, v)$, 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 the convexity condition stated in [2] implies convexity of U_v. If the arbitration function is refined, we provide a characterization of coalition structures that can be stabilized which is equivalent to the Bondareva-Shapley [1, 5] characterization of stable non-OCF games. This characterization allows us to identify a sufficient condition on v which ensures that the refined core is not empty: if the superadditive cover of v is convex, then the refined core of (N, v) 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. This essentially means that LBGs can be stabilized regadless of agent reaction to deviation, since the optimistic core is contained in all other arbitrated cores.

5. MANIPULATING THE QUOTA IN WEIGHTED VOTING GAMES

In a *weighted voting game (WVG)*, there is an integer weight w_i assigned to each agent and a quota q. Given a set of agents S, we define $w(S) = \sum_{i \in S} w_i$; $w(S)$ is called the weight of S. S has value 1 if and only if $w(S) \geq q$, and is 0 otherwise. This model is useful in describing the behavior of decision making agents. It is often useful to measure the relative power of individual agents, i.e. how central is an agent in the decision making process. One such measure of influence is the Shapley value of a player i, denoted φ_i; the higher the value, the more influential the agent. In [9], we study the effects of changes to the quota on the Shapley value of players in a WVG; thus, φ_i becomes a function of q, denoted $\varphi_i(q)$.

In [9] we state several decision problems:

1. Is there an efficient method of finding a quota q_0 for which φ_i is maximal/minimal
2. Is there an efficient method of deciding whether the value of $\varphi_i(q)$ is maximal/minimal?

First, we show that finding a maximizing q_0 is easy: simply set $q_0 = w_i$. Second, we show that both versions of (2) are NP-hard. Finally, we show via empirical analysis that the answer to the minimization version of (1) is usually setting the quota to either 1 or $w_i + 1$, with the preferred setting depending on the agent's relative weight; in particular, we show that if an agent's weight is below the median weight, it is always preferable to set the quota to $w_i + 1$ if one wishes to minimize φ_i.

6. REFERENCES

[1] O. Bondareva. Some applications of linear programming methods to the theory of cooperative games. *Problemy kibernetiki*, 10:119–139, 1963.

[2] G. Chalkiadakis, E. Elkind, E. Markakis, M. Polukarov, and N. Jennings. Cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research*, 39:179–216, 2010.

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

[4] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*, volume 34 of *Theory and Decision Library. Series C: Game Theory, Mathematical Programming and Operations Research*. Springer, Berlin, second edition, 2007.

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

[6] Y. Zick, G. Chalkiadakis, and E. Elkind. Overlapping coalition formation games: Charting the tractability frontier. In *Proc. of the 11th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 2012.

[7] Y. Zick and E. Elkind. Arbitrators in overlapping coalition formation games. In *Proc. of the 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 55–62, 2011.

[8] Y. Zick, E. Markakis, and E. Elkind. Stability via LP duality in games with overlapping coalitions. Submitted to the 26th Conference on Artificial Intelligence (AAAI 2012), 2012.

[9] Y. Zick, A. Skopalik, and E. Elkind. The Shapley value as a function of the quota in weighted voting games. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 490–495, 2011.