Artificial Intelligence 232 (2016) 76–113

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Characteristic function games with restricted agent interactions: Core-stability and coalition structures

Georgios Chalkiadakis [a], Gianluigi Greco [b,*], Evangelos Markakis [c]

[a] Technical University of Crete, Chania, Greece
[b] Department of Mathematics and Computer Science, University of Calabria, Italy
[c] Athens University of Economics and Business, Athens, Greece

ARTICLE INFO

Article history:
Received 15 January 2014
Received in revised form 14 December 2015
Accepted 17 December 2015
Available online 23 December 2015

Keywords:
Coalitional games
Solution concepts
Computational complexity
Treewidth
Marginal contribution networks

ABSTRACT

In many real-world settings, the structure of the environment constrains the formation of coalitions among agents. These settings can be represented by characteristic function games, also known as *coalitional games*, equipped with interaction graphs. An interaction graph determines the set of all feasible coalitions, in that a coalition C can form only if the subgraph induced over the nodes/agents in C is connected. Our work analyzes stability issues arising in such environments, by focusing on the *core* as a solution concept, and by considering the coalition structure viewpoint, that is, without assuming that the grand-coalition necessarily forms.
The complexity of the coalition structure core is studied over a variety of interaction graph structures of interest, including complete graphs, lines, cycles, trees, and nearly-acyclic graphs (formally, having bounded treewidth). The related stability concepts of the *least core* and the *cost of stability* are also studied. Results are derived for the setting of *compact* coalitional games, i.e., for games that are implicitly described via a compact encoding, and where simple calculations on this encoding are to be performed in order to compute the payoff associated with any coalition. Moreover, specific results are provided for compact games defined via *marginal contribution networks*, an expressive encoding mechanism that received considerable attention in the last few years.

1. Introduction

1.1. Coalitional games and interaction graphs

Cooperative game theory aims to provide models of cooperation among interacting agents. One of the prevalent classes of games used within this framework is the class of *characteristic function games*. A characteristic function game (CFG) is defined over a set $N = \{1, \ldots, n\}$ of agents and is determined by a *payoff function* $v : 2^N \mapsto \mathbb{R}$, such that, for each *coalition* C, i.e., for any non-empty set $C \subseteq N$ of agents, the value $v(C)$ expresses the payoff that the members of C can jointly achieve by cooperating among themselves [1,2]. The outcome is an *allocation*, i.e., a payoff vector $\boldsymbol{x} = \langle x_1, \ldots, x_n \rangle \in \mathbb{R}^n$ assigning some payoff to each agent $i \in N$. Characteristic function games are also known as *coalitional games with transferable utility*, as it is assumed that the agents forming a coalition C can distribute the payoff $v(C)$ among themselves in any way. The

* Corresponding author.
E-mail addresses: gehalk@intelligence.tuc.gr (G. Chalkiadakis), ggreco@mat.unical.it (G. Greco), markakis@gmail.com (E. Markakis).

http://dx.doi.org/10.1016/j.artint.2015.12.005

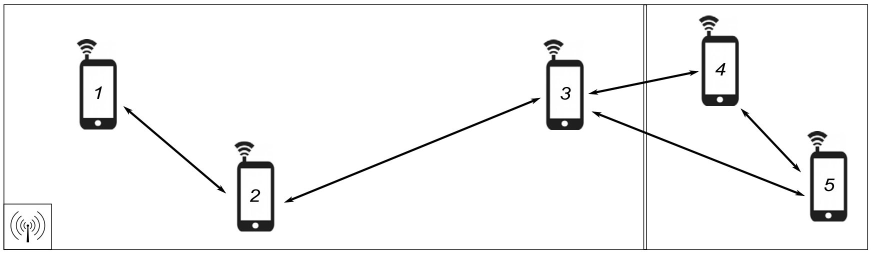

Fig. 1. Coalitional games in Example 1.1 and Example 1.3.

question of interest for such games, is to identify desirable (e.g., fair and stable) outcomes in terms of worth distributions, which are called *solution concepts*.

Characteristic function games provide a rich framework for understanding and reasoning about cooperative actions and have a wide spectrum of applications in different areas of research. Indeed, while they have been traditionally grounded in the game theory and economic literature, they have gained popularity in the context of multi-agent systems and artificial intelligence research as a means of studying interactions among autonomous agents (see, e.g., [3–9]). Moreover, they have recently attracted attention in engineering too, because of their use in the design of intelligent protocols and middleware algorithms for wireless communication networks [10–12]. As an example, we next illustrate a *wireless cooperative file sharing system*, where mobile subscribers cluster together by downloading (portions of) files of interest over long-range cellular links, and by exchanging them over a wireless ad-hoc network (see, e.g., [13]).

Example 1.1. Consider the setting with three mobile users (hence, $N = \{1, 2, 3\}$) that is illustrated in the left part of Fig. 1, and a function v such that, for each coalition $C \subseteq N$, $v(C)$ is meant to express the payoff that users in C can jointly achieve when clustering together and cooperating to download the given file of interest. In particular, assume that $v(\{1\}) = 10$, $v(\{2\}) = 5$, $v(\{3\}) = 4$, $v(\{1, 2\}) = 17$, $v(\{2, 3\}) = 10$, and $v(\{1, 2, 3\}) = 22$. Intuitively, the payoff of each coalition C is meant to express the sum of the utilities that the agents in C get when the file is downloaded minus the cost that they overall incur for the download, with this cost being proportional to the time required and to certain technological features, such as bandwidth and energy consumption.

For instance, when 1 and 2 cluster together, each of them can download only half of the file and then share the portion via the wireless connection. While doing so, each of them increases the throughput due to the better performance of the wireless network when compared to the cellular links, and reduces the overall downloading costs. This is reflected in the payoff function v, which is such that $v(\{1, 2\}) \geq v(\{1\}) + v(\{2\})$. Similarly, 2 and 3 can cluster together leading to the payoff $v(\{2, 3\}) \geq v(\{2\}) + v(\{3\})$.

In contrast, note that users 1 and 3 are outside of each other's transmission range, hence it is impossible for them to cluster together (and, in fact, the payoff function is not specified for this coalition). However, user 2 might still act as a bridge between them, so that when all users join together, the resulting payoff is $v(\{1, 2, 3\}) = 22$. Indeed, it is advantageous for the users to cluster all together, because $v(\{1, 2, 3\}) \geq v(\{1, 2\}) + v(\{3\})$, $v(\{1, 2, 3\}) \geq v(\{1\}) + v(\{2, 3\})$, and $v(\{1, 2, 3\}) \geq v(\{1\}) + v(\{2\}) + v(\{3\})$. While doing so, each user gets the payoff that can be achieved by downloading the whole file alone, and a surplus of $v(\{1, 2, 3\}) - (v(\{1\}) + v(\{2\}) + v(\{3\})) = 22 - (10 + 5 + 4) = 3$ still remains to be divided among them. $\lhd$

As the above example demonstrates, characteristic function games might be defined within an environment imposing restrictions on the formation of coalitions. Indeed, users 1 and 3 are outside each other's transmission range, and the coalition $\{1, 3\}$ cannot form.

In general, for reasons that might range from physical limitations and constraints to legal banishments, certain agents might not be allowed to form coalitions with certain others. Sensor networks, communication networks, or transportation networks, within which units are connected through bilateral links, provide natural settings for such classes of games. In many multiagent coordination settings, agents might be restricted to communicate or interact with only a subset of other agents in the environment, due to limited resources or existing physical barriers. Another example is provided by hierarchies of employees within an enterprise, where employees at the same level work together under the supervision of a boss, i.e., of an employee at the immediately higher level in the hierarchy. In all these settings, the environment can be seen to possess some *structure* that forbids the formation of certain coalitions. This can be formalized as an *interaction graph* $G = (N, E)$, an undirected graph, where agents are transparently viewed as nodes so that a coalition C is *feasible*, i.e., allowed to form, only if the subgraph of G induced over the nodes of C is connected [14]. For instance, it is immediate to check that the graph shown in the left part of Fig. 1, is the interaction graph associated with the game of Example 1.1, where all coalitions are allowed to form but $\{1, 3\}$.

Note that when an interaction graph G is the *complete graph* over N, then G induces no structural restrictions and we are back to the basic setting where a game is completely specified by its payoff function (and all coalitions are allowed to form). Hence, the setting with interaction graphs generalizes the basic one.