

Coalitional Stability in Structured Environments

Georgios Chalkiadakis
Technical University of Crete
Chania, Greece
gehalk@ece.tuc.gr

Evangelos Markakis
Athens University of
Economics and Business
Athens, Greece
markakis@gmail.com

Nicholas R. Jennings
University of Southampton
Southampton, UK
nrj@ecs.soton.ac.uk

ABSTRACT

In many real-world settings, the structure of the environment constrains the formation of coalitions among agents. Therefore, examining the stability of formed coalition structures in such settings is of natural interest. We address this by considering core-stability within various models of cooperative games with structure. First, we focus on characteristic function games defined on graphs that determine feasible coalitions. In particular, a coalition S can emerge only if S is a connected set in the graph. We study the (now modified) core, in which it suffices to check only feasible deviations. Specifically, we investigate core non-emptiness as well as the complexity of computing stable configurations. We then move on to the more general class of (graph-restricted) partition function games, where the value of a coalition depends on which other coalitions are present, and provide the first stability results in this domain. Finally, we propose a “Bayesian” extension of partition function games, in which information regarding the success of a deviation is provided in the form of a probability distribution describing the possible reactions of non-deviating agents, and provide the first core-stability results in this model also.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Economics

Keywords

cooperative game theory, coalition formation, core

1. INTRODUCTION

Cooperative game theory, providing as it does a rich framework for the study of coalition formation among rational players, has in recent years attracted much attention in multiagent systems as a means of forming teams of autonomous agents. The vast majority of work in cooperative game theory assumes that, given a set of agents, any coalition among them is allowed to form. However, in many circumstances the environment imposes restrictions on the formation of coalitions: for reasons that might range from physical limitations and constraints to legal banishments, certain agents

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

might not be allowed to form coalitions with certain others. 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. In such settings, the environment can be seen to possess some *structure* that forbids the formation of certain coalitions. This can be captured by an undirected graph providing a path connecting any two agents that can belong to the same coalition.

Specifically, sensor networks, communication networks, or transportation networks, within which units are connected through bilateral links, provide natural settings for cooperative games defined over graphs. Another example is provided by hierarchies within an enterprise, where the underlying graph corresponds to a tree.

In this paper we consider various models of cooperative games in *structured* environments as above and study the *stability* of coalition structures in such settings. Stability is one of the key issues in cooperative game settings and examines whether agents have an incentive to depart from an existing coalition structure. As an example, airlines participating in a certain alliance may be willing to move to an alliance that could guarantee them higher profits. Here we focus on the celebrated cooperative stability concept of *the core* [21], which is the set of outcomes that are stable against deviations by any subset of agents. In our work, the definition of the core has to be modified so that the only allowed deviations are those by sets of *connected* coalitions.

Against this background, the rest of the paper is structured as follows. We start with the usual *characteristic function games* (CFGs) setting, under the assumption, however, that the games are defined over graphs. We introduce to the community recent results from the economics literature, which establish the non-emptiness of the core in games defined over a tree, and determine a procedure to find a core element in such games [5]. We then focus on three natural graph structures—lines, trees and cycles—and study the computational problems of (i) deciding the non-emptiness of the core; (ii) finding an element in the core; and (iii) checking if a given outcome belongs to the core. We show certain *positive* results when the underlying graph is a line or a cycle, and when it is a tree and the game is superadditive. These results are interesting from an applications point of view, since many computer or sensor networks exhibit a ring or tree topology. However, for non-superadditive games over trees, certain *negative* complexity results are obtained.

Then, we move on to the more general class of *partition function games* (PFGs) over graphs, and initiate the study of stability in that setting. In PFGs, the value of a coalition depends on the partition currently in place [25]. Defining the core in the presence of *externalities* is complicated and there is no unanimously accepted solution as potential deviators in PFGs have to consider how non-deviators—the “residual” players—would react to their deviation.

Since residual players can form any structure among themselves, the value of any deviation relies on the resulting partition across the space of agents. A common treatment in the literature is for the deviators to either *pessimistically* assume that non-deviators will partition so as to hurt them the most, or to *optimistically* assume that the partition of the non-deviators will be the best possible [23, 8]. We adopt both those views in turn, and provide the first core non-emptiness results in PFG settings with structure. Operating first under the assumption of *pessimism*, we define the (pessimistic) core and then show that for any PFG, there is a corresponding CFG, such that the core of the CFG is contained in the core of the PFG (but the opposite is not always true). Interestingly, this differs from what is known to hold for the class of PFGs where the coalition of all agents is the partition with maximum social welfare (in which case the two cores coincide [8]). This correspondence enables us to generalize the CFG-related results, and show that the core is non-empty for PFGs defined over trees. Furthermore, the same process as before can be used to obtain a core-stable configuration in a PFG defined over a tree, however this will not generally run in polynomial time. We then adopt an *optimistic* view regarding the behaviour of non-deviators, and show that, unlike the pessimistic one, the optimistic core may be empty in PFGs over trees and even over lines. It remains an interesting future work topic to obtain efficient algorithms for special cases of PFGs. As explained in Section 3.1, this seems to be computationally much harder, since it typically involves enumeration over too many partitions. We are not aware of any other work that has tackled stability in PFG settings with structure.

Finally, we propose a natural extension of PFGs, namely *Bayesian partition function games (BPFGs)*. In short, instead of resorting to pessimism or optimism, a coalition S of potential deviators in BPFGs assumes that the reaction of the residual players (i.e., what partition they will form if S deviates) is determined by a probability distribution—an assumption that is more realistic and arguably more useful from an AI perspective. We then go on to define the core and initiate its study in this setting as well.

2. CHARACTERISTIC FUNCTION GAMES ON GRAPHS

In this section, we study the issue of stability in the context of characteristic function games defined on graphs. Let $N = \{1, \dots, n\}$ be a set of agents, with $|N| = n$. A subset $C \subseteq N$ is called a coalition. A *characteristic function game (CFG)*—or *coalitional game with transferable utility (TU-game)*—is defined by its function $v : 2^N \mapsto \mathfrak{R}$ that specifies the value $v(C)$ of each coalition C [21]. Intuitively, $v(C)$ represents the maximal payoff the members of C can jointly receive by cooperating, and the agents can distribute this payoff between themselves in any way. A payoff vector $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ assigning some payoff to each $i \in N$ is called an *allocation*. We denote $\sum_{i \in C} x_i$ by $x(C)$. Given a partition $\Pi = \{C_1, \dots, C_k\}$, of the agents (we will also refer to a partition as a *coalition structure* interchangeably), an allocation \mathbf{x} is called an *imputation* of Π if $x(C_j) = v(C_j)$ for $j = 1, \dots, k$ and $x_i \geq v(\{i\})$ for all i . Note that if \mathbf{x} is an imputation for Π , $\sum_{i \in N} x_i = \sum_{C \in \Pi} v(C)$. The set of imputations for Π is $I(\Pi)$.

Assume now that there exists a graph G , which determines the allowed cooperation structures as follows: each node of the graph represents an agent and a coalition C is allowed to form if and only if for every two agents in C there exists a path in the subgraph induced by C that connects them—i.e., the subgraph that is induced by C is a connected subgraph. A *characteristic function game on graph G* is then simply a CFG where v is defined only for coalitions allowed by G . We denote the set of such *feasible coalitions*

by $\mathcal{F}(G)$. Similarly, we will refer to *feasible partitions* of feasible coalitions, and we denote the set of all feasible partitions by $P(G)$. Such games can arise naturally in many situations where lack of communication between certain agents makes it impossible for some coalitions to form. Note that when G is a clique then we are back to the usual CFGs where all coalitions are feasible. We also assume that our graph is connected. If not, our findings apply separately to each connected component.

The main stability solution concept in cooperative game theory is, arguably, *the core*—the set of (Π, \mathbf{x}) tuples, where Π is a partition and \mathbf{x} an imputation, such that no coalition has an incentive to deviate. However, in our setting it suffices to check only the incentives of the feasible coalitions [17, 4, 5].

DEFINITION 1. *The core of a game with characteristic function $v(\cdot)$ on graph G is the set*

$$C(v, G) = \{(\Pi, \mathbf{x}) : \Pi \in P(G), \mathbf{x} \in I(\Pi) \wedge x(S) \geq v(S) \forall S \in \mathcal{F}(G)\}$$

The following observation is straightforward:

FACT 1. *If $(\Pi, \mathbf{x}) \in \text{core}$, Π attains maximum social welfare, where the social welfare of Π is: $SW(\Pi) = \sum_{C \in \Pi} v(C)$.*

Though in many games the *grand coalition* of all players might be impossible to form [1], the assumption that it is the one with the highest total welfare—or even the stricter assumption of *superadditivity*, i.e., $v(S \cup T) \geq v(S) + v(T)$ for any disjoint sets S, T —is some times justified. Indeed, the vast majority of work in game theory examines stability in games where the grand coalition emerges. The question of stability then reduces to pairing the grand coalition with an imputation of N .¹

DEFINITION 2. *When N attains the highest possible social welfare, the core is the set*

$$C(v, G) = \{\mathbf{x} \in \mathfrak{R}^n : x(N) = v(N) \text{ and } x(S) \geq v(S) \forall S \in \mathcal{F}(G)\}$$

In most scenarios of interest to multiagent systems, however, it is natural for agents to split into groups to simultaneously perform distinct tasks. Thus, unless explicitly stated, it is *not* required in our games that the grand coalition achieves the highest social welfare.

Network structures have long been recognized as a natural framework for the study of stability. Nevertheless, following the definition of *Myerson value*² in [20], research has focused on the question of *building* stable and efficient networks: (mainly pairwise) stability is discussed essentially from a non-cooperative point of view—i.e., w.r.t. *the creation* of stable network structures, through adding or removing links among nodes [14]. In non-cooperative settings, the structural properties of equilibria and the development of algorithms to compute equilibria in *graphical games* which restrict payoff influences among players have also been examined [15].

Networks have also provided inspiration for new representation schemes for coalitional games [13, 3, 2]. Cooperative games with an underlying graph structure have also been considered in the seminal work of [6]. However, they consider games defined on a weighted graph, where the value of a coalition S is the sum of weights of edges that are contained in the subgraph induced by S . Hence, *any* coalition is allowed to form and values are determined by weights, while in our work not all coalitions are feasible and

¹The core of Definition 1, examining the stability of coalition structures is sometimes referred to as *the CS-core*, while the concept examining the stability of N is referred to simply as *the core* [7].

²The Myerson value determines that, while building a network by adding agents one at a time, each agent is assigned his marginal contribution, taking into account all possible orderings of agents.

the characteristic function can be arbitrary and not expressed via weights. Thus, their results are unrelated to ours.

In contrast to the aforementioned approaches, here we care about the stability of *partition-imputation pairs* given a *fixed* underlying network structure that determines the allowed interactions. In particular, we study the complexity of three³ natural core-related algorithmic questions (Table 1 summarizes our results):

1. **CORE-NONEMPTINESS:** Given a game on a graph G , decide whether $C(v, G) \neq \emptyset$.
2. **CORE-FIND:** Given a game on a graph G , find an element $(\Pi, x) \in C(v, G)$ if $C(v, G) \neq \emptyset$ or output “ $C(v, G) = \emptyset$ ”.
3. **CORE-MEMBERSHIP:** Given a game on a graph G and an imputation (Π, x) , decide whether $(\Pi, x) \in C(v, G)$.

REMARK 1. *We make the usual assumption that our games are in compact form and are represented by the graph G and an oracle that, for any $S \in \mathcal{F}(G)$, returns the value $v(S)$ in time polynomial in the size of G (e.g. see [9]).*

	Lines (general)	Trees (superadd.)	Trees (general)	Cycles (general)
NONEMPTINESS	O(1)	O(1)	O(1)	P
FIND	P	P	NP-hard	P
MEMBERSHIP	P	co-NP-complete	co-NP-complete	P

Table 1: Core-stability results for CFGs on graphs.

2.1 CORE-NONEMPTINESS and CORE-FIND

We start with the problems of determining whether the core is empty or not and the complexity of finding elements in the core. Related work on solution concepts in graph-restricted games [18, 26] has not addressed issues from an algorithmic point of view, in most cases. The work most relevant to ours is that of Le Breton *et al.* [17], and Demange [4, 5]. In [17] and [4] it is shown that if a game is superadditive and the graph is a tree, the core is non-empty. However, their existential proof does not provide an efficient algorithmic construction of a core element.

Later on, the follow up work of [5] showed that the core is non-empty for trees, even for non-superadditive games. Moreover, Demange proposed a procedure that computes an element in the core. We briefly recall this algorithm below as we will proceed to analyze its complexity, and will also use it in later sections. Originally it was stated in the context of a slightly different model than ours, involving directed graphs, but it is easy to reformulate as:

ALGORITHM 1. [5] *Given a graph G which is a tree, and a characteristic function $v(\cdot)$, first pick a vertex r as the root of the tree. The algorithm consists of two steps.*

Step 1: *Starting from the leaves, compute the guarantee level \hat{g}_i of each agent i , inductively: for leaves, \hat{g}_i is simply the reservation value $v(\{i\})$; for an agent that is not a leaf, let \mathcal{R}^i be the set of all subtrees that start at i , i.e., it is the set of all feasible coalitions among i and the agents underneath i . Then: $\hat{g}_i := \max\{v(T) - \sum_{j \in T \setminus \{i\}} \hat{g}_j : T \in \mathcal{R}^i\}$.*

Step 2: *Starting from the root r , pick the coalition T_1 at which \hat{g}_r was attained (breaking ties arbitrarily). Every agent in T_1 receives his guarantee level as a payoff, which is feasible by the definition of \hat{g}_r . If $T_1 = N$, we are done. Otherwise, pick a node i not in T_1 whose father belongs to T_1 . Pick the coalition from \mathcal{R}^i at which \hat{g}_i*

³The first problem is included in the second one. However, we feel it is important to study CORE-NONEMPTINESS separately from CORE-FIND because their complexity varies significantly.

was attained, say T_2 . All agents in T_2 receive their guarantee level as well. If $T_1 \cup T_2 = N$, we are done, otherwise we continue in the same fashion until we cover N . This produces a partition Π in which all agents receive their guarantee level as their payoff.

THEOREM 1 ([5]). *The outcome produced by the above algorithm belongs to the core.*

To obtain more intuition, it is interesting to observe what the algorithm does in some special cases:

REMARK 2. *If the game is superadditive and the graph is a line from 1 to n , the produced payoff allocation is simply the marginal contribution: $\hat{g}_i = v(\{1, \dots, i\}) - v(\{1, \dots, i-1\})$.*

The above theorem implies that CORE-NONEMPTINESS is trivial when the graph is a tree (whether superadditive or not), since the core is always non-empty. Regarding CORE-FIND however, the computational complexity of Algorithm 1 was not addressed in [5]. We therefore now proceed to analyze its complexity. In the usual CFGs, where there is no restriction by a graph, the problem of computing an element in the core, or deciding if the core is non-empty has already been shown to be in co-NP, and co-NP-complete for certain expressive representation schemes [11]. Here, the fact that the graph may restrict the number of potential deviations gives some hope that the problem may be easier. We show below that this is indeed the case for graphs that are lines or superadditive trees.⁴

THEOREM 2. *For a CFG where (i) the underlying graph is a line or (ii) the game is superadditive and the graph is a tree, CORE-FIND can be solved in polynomial time.*

PROOF. (i) For finding an element in the core it is easy to see that Algorithm 1 runs in polynomial time for lines. To prove this, we need to see how many coalitions we need to check when we compute the \hat{g}_i values for each agent i . Suppose wlog that the agents are placed in a line starting from agent 1 up to agent n . For agent 1, since it is a leaf, the only coalition we consider is the singleton $\{1\}$. For player 2, to compute \hat{g}_2 , we need to consider $\{2\}$ and $\{1, 2\}$. Moving on this way we see that for agent $n-1$ we need to consider $\{n-1\}, \{n-2, n-1\}, \dots, \{1, \dots, n-1\}$. And finally for agent n , the allowed coalitions are $\{n\}, \{n-1, n\}, \dots, \{1, \dots, n\}$. In total for all players we need to consider $1 + 2 + \dots + (n-1) + n = O(n^2)$ coalitions. Hence it is a polynomial time algorithm.

(ii) Suppose now that the game is superadditive and the graph is a tree. The computational problem that may arise on a tree is the following: if a node i has k children then computing its guarantee level in step 1 of Algorithm 1 requires looking at exponentially many subtrees starting at i , i.e. the set \mathcal{R}^i contains at least 2^k subtrees corresponding to all possible subsets of the children. However, when the game is superadditive, this is not necessary, as implied by the following:

LEMMA 1. *For a node i , let D_i be the tree that starts at i and contains all nodes downwards from i . When the game is superadditive, the guarantee level of every node in step 1 of Algorithm 1 is achieved precisely at D_i .*

The proof of Lemma 1 follows by induction and we omit it here. Given Lemma 1, we can conclude that Algorithm 1 can be implemented in polynomial time in this case. \square

⁴Notice, however, that *testing* whether a game is indeed superadditive could still be a hard problem—e.g., it is coNP-complete in the absence of structured environments [10].

Given the results of the above theorem, and Fact 1, it is straightforward to obtain the following corollary:

COROLLARY 1. *For a CFG where the underlying graph is a line, or it is superadditive and the graph is a tree, a partition with maximum social welfare can be found in poly-time.*

When the game is not superadditive, however, CORE-FIND becomes intractable for trees. To see this, we first define:

The social welfare maximization (SW) problem: *Given a game on a graph G , and a rational number k , is there a partition Π with $SW(\Pi) \geq k$?*

THEOREM 3. *The SW problem when the underlying graph is a tree is NP-complete.*

PROOF. That SW is in NP is trivial. Just guess a partition and calculate its social welfare. To show that it is NP-hard, we reduce from the PARTITION problem, which is the following: given n positive numbers a_1, \dots, a_n is there a subset S of these numbers such that $\sum_{j \in S} a_j = \sum_{j \notin S} a_j$?

Consider an arbitrary instance of the PARTITION problem with numbers a_1, \dots, a_n . We construct a CFG on a graph which is a tree. In particular, we define a graph with $n + 1$ vertices, namely the vertices $\{0, 1, \dots, n\}$. The only edges are the edges $(0, i)$ for $i = 1, \dots, n$. Hence the graph is a tree rooted at vertex 0. To determine the game's characteristic function, notice first that feasible coalitions of size at least 2 are forced to contain the root, otherwise they are not connected. Therefore the only allowable coalitions are singletons and sets that contain vertex 0. For the singleton $\{0\}$, set $v(\{0\}) = 0$. For the rest of the singletons, set $v(\{i\}) = a_i$. Finally, the value of coalitions S that contain the root is:

$$v(S) = \begin{cases} \sum_{j \in S \setminus \{0\}} a_j & \text{if } \sum_{j \in S \setminus \{0\}} a_j = \sum_{j \notin S} a_j, \\ 1 + \sum_{j \in S \setminus \{0\}} a_j & \text{otherwise} \end{cases}$$

By the definition of $v(S)$, it is easily seen that given a_1, \dots, a_n , a polynomial time oracle for $v(S)$ can be constructed. Set now $k = 1 + \sum_{j=1}^n a_j$. This completes the description of the SW instance. We can now prove that there exists a set S such that $\sum_{j \in S} a_j = \sum_{j \notin S} a_j$ in the PARTITION problem if and only if there exists a feasible partition with social welfare at least k . Suppose there exists such a set S . Then for the coalition $S \cup \{0\}$ the value is exactly $1 + \sum_{j \in S \setminus \{0\}} a_j$. Hence with the remaining nodes as singletons, we get a partition with social welfare at least k . For the reverse, suppose that there exists a partition with social welfare at least $1 + \sum_{j=1}^n a_j$. If there is no set S that solves the PARTITION problem, then by construction the values of all coalitions are just the sum of the corresponding numbers, and hence the social welfare of any feasible partition is at most $\sum_{j=1}^n a_j$, a contradiction. \square

Now, we can see why our problem of finding an element in the core is unlikely to have a polynomial time algorithm⁵:

THEOREM 4. *For a general CFG where the underlying graph is a tree, CORE-FIND is NP-hard.*

PROOF. Given Fact 1, any algorithm that finds an element (Π, \mathbf{x}) in the core can solve the SW problem, by just calculating the social welfare of the returned Π . \square

We now focus on games where the underlying graph is a cycle. These are interesting, as ring topologies are common in networks

⁵Since our problem is a search problem, it will not belong to NP. Our proof essentially shows that our problem is at least as hard as solving an NP-complete problem.

of many kinds. It can be shown through examples that trees (and forests, if G is disconnected) are the only graphs that guarantee stability irrespective of the function $v(\cdot)$. The presence of cycles can create instances with empty core, even in superadditive games.

THEOREM 5. *For any $n \geq 3$, there exist games on n players, where the underlying graph is a cycle and $C(v, G) = \emptyset$.*

For $n = 3$, Theorem 5 follows by the abundance of regular unrestricted CFG's with empty core (since a cycle is a clique for $n = 3$, and hence all coalitions are allowed). For $n \geq 4$ we can construct simple examples of superadditive characteristic functions on cycles, which we omit here.

Hence, the problem CORE-NONEMPTINESS for cycles is not as trivial as in the case of trees. We will show however that we can still have a polynomial time algorithm. We start with superadditive games on cycles. The crucial observation is the following Lemma:

LEMMA 2. *For lines and cycles, the number of feasible coalitions is $O(n^2)$.*

PROOF. A feasible coalition has to be connected and it corresponds to an interval from an agent i to some agent j . This implies a total of $O(n^2)$ since we have at most n choices for i and after fixing i , there can be at most $n - 1$ choices for j (multiplied by 2 for cycles, since we then have two paths connecting i and j). \square

THEOREM 6. *CORE-NONEMPTINESS and CORE-FIND are in P for superadditive games on cycles.*

PROOF. From superadditivity, we know that we are looking for an imputation \mathbf{x} of the grand coalition. Hence we can check if the following system of linear inequalities has a solution:

$$\sum_{i \in N} x_i = v(N) \text{ and } \sum_{i \in S} x_i \geq v(S) \quad \forall S \in \mathcal{F}(G)$$

By Lemma 2, we know that this system has polynomially many constraints and hence can be solved in polynomial time. \square

Note that we could have applied the same argument for superadditive games on lines. However, for lines we prefer to use Algorithm 1, since it works for non-superadditive games as well, and it provides a more direct and intuitive way of finding a core element.

One cannot directly use the arguments in Theorem 6 for non-superadditive games on cycles, as an optimal partition is not known *a priori*, and there are exponentially many candidate partitions. However, if one has access to an algorithm that computes an optimal partition Π , then an allocation vector \mathbf{x} so that (Π, \mathbf{x}) is in the core, if such an \mathbf{x} exists, can be computed via linear programming, by the arguments above. Hence, we need to address the question of whether an optimal partition can be computed in poly-time *over cycles*. Let C be a cycle over n nodes, and let L_1, \dots, L_n be the n lines that can be obtained by removing exactly one edge. Then, it is easy to see that the optimal partition of C is either C itself, or the best partition over the optimal partitions of L_1, \dots, L_n . Thus, one can just run Alg. 1 on L_1, \dots, L_n and compare the various solutions with the value of the grand coalition. This proves that CORE-NONEMPTINESS and CORE-FIND are in P for general games on cycles. Credit for this proof (provided to us in personal correspondence) goes to G. Greco, E. Malizia, L. Palopoli and F. Scarcello.

Finally, we conclude this subsection with an observation on how to identify more core elements if one has access to algorithms that identify alternative optimal partitions (a result also shown in [10]).

THEOREM 7. *For any game on G , if $(\Pi, \mathbf{x}) \in C(v, G)$, then $(\Pi', \mathbf{x}) \in C(v, G)$, for any social welfare maximizing Π' .*

2.2 CORE-MEMBERSHIP

In this section we deal with the membership problem. First we show that it can be solved efficiently for lines or cycles.

THEOREM 8. *For a CFG where the underlying graph is a line or a cycle, CORE-MEMBERSHIP can be resolved in \mathbf{P} .*

PROOF. Given (Π, \mathbf{x}) , it is trivial to check if x is an imputation of Π . To check for a successful deviation, we can check all feasible coalitions. By Lemma 2 the proof is complete. \square

Concerning trees, the problem is not as easy. We first show that the reduction we used in the proof of Theorem 3 yields a hardness result for general games on trees.

THEOREM 9. *CORE-MEMBERSHIP is co-NP-complete for general CFGs on trees.*

PROOF. To check that a (Π, \mathbf{x}) does not belong to the core, it suffices to exhibit a coalition with an incentive to deviate, or check that \mathbf{x} is not an imputation of Π . Hence there is a polynomially sized certificate to verify this, which implies that CORE-MEMBERSHIP belongs to co-NP (also follows from a result in [10]).

We now show that checking whether (Π, \mathbf{x}) does not belong to the core is NP-hard. We use the reduction from the proof of Theorem 3. Given an instance of the PARTITION problem (a_1, \dots, a_n) , we construct the game described in Theorem 3 and we consider as a candidate core element the tuple $(\Pi, \mathbf{x}) = (N, (0, a_1, \dots, a_n))$ —i.e., we look at the grand coalition with the imputation where the root receives nothing and every other node receives its corresponding number. We claim that (Π, \mathbf{x}) does not belong to the core if and only if there exists a set S that is a solution to the PARTITION problem. To see this, suppose first that there exists a set S such that $\sum_{j \in S} a_j = \sum_{j \notin S} a_j$. Then, by construction and the definition of the characteristic function, the grand coalition is not welfare maximizing, and thus cannot belong to the core (there exists a partition with social welfare at least $1 + \sum_{j=1}^n a_j$).

For the other direction suppose that there is no set S for which $\sum_{j \in S} a_j = \sum_{j \notin S} a_j$. This implies that the grand coalition is a welfare maximizing partition and it is also easy to see that every coalition achieves its value $v(S)$ under the imputation \mathbf{x} . Hence (Π, \mathbf{x}) belongs to the core. \square

The reduction above does not imply anything for superadditive trees since it produces non-superadditive instances as well. One could hope that superadditivity makes things easier as was the case for CORE-FIND. The following theorem however reveals that CORE-MEMBERSHIP is hard even for superadditive trees. The proof is based on a simple adaptation of a reduction used for unrestricted games in [10][Theorem 4.1].

THEOREM 10. *CORE-MEMBERSHIP is co-NP-complete even for superadditive CFGs on trees.*

PROOF. (Sketch.) Checking that the problem is in co-NP is as in Theorem 9. We now give a reduction from SAT. Consider a SAT formula ϕ on n boolean variables X_1, \dots, X_n . Let L denote the set of literals, $L = \{X_1, \neg X_1, X_2, \neg X_2, \dots\}$. Given $S \subseteq L$, with $|S| = n$, we say that S is consistent if $\forall X_i \ |S \cap \{X_i, \neg X_i\}| = 1$. For a consistent S , let also $\sigma(S)$ be the truth assignment where all variables that belong to S are set to true and the rest to false.

Given ϕ , we construct a tree with $2n + 1$ nodes. The root is denoted by node 0 and it is connected to $2n$ nodes corresponding to L . This is a star where the only allowable coalitions are either singletons or coalitions containing the root. For any singleton coalition $\{i\}$ we set $v(\{i\}) = 0$. All remaining coalitions are of the form

$\{0\} \cup S$ for some $S \subseteq L$. The value of $v(\{0\} \cup S)$ is set to the value of the set S in the proof of [10][Theorem 4.1]. Namely:

$$v(\{0\} \cup S) = \begin{cases} |S|/n, & \text{if } |S| > n, \\ 1 + 1/2n & \text{if } |S| = n, S \text{ is consistent,} \\ & \text{and } \sigma(S) \text{ satisfies } \phi \\ 0, & \text{otherwise} \end{cases}$$

It can be easily proved that this game is superadditive. Furthermore, one can also prove that the allocation where node 0 receives 0 and every other node $1/n$ belongs to the core if and only if the formula ϕ is not satisfiable. We omit the details from this version. \square

3. PARTITION FUNCTION GAMES

As mentioned in the introduction, in many circumstances the value of a coalition S does not depend solely on S but is affected by externalities—i.e., $v(S)$ depends on the coalition structure formed by the rest of the agents. To capture such requirements, one is then obliged to move to the more general setting of partition function games. Naturally, therefore, it is of interest to extend the results of Section 2 to *partition function games with structure*. As discussed, the core of PFGs has been studied in economics under specific assumptions. Moreover, recent work has focused on efficient PFG representations and coalition structure generation in PFG settings (Michalak *et al.* [19]; Rahwan *et al.* [22]). However, to date there has been no work on the stability of PFGs defined on graphs.

We begin with some definitions. A PFG is determined by a function $V(\cdot, \cdot)$, specifying the value of a coalition in a certain partition. For $S \in \Pi$, $V(S, \Pi)$ is the value of S when Π forms. In our setting, to define a PFG on a graph G , we further impose that the function $V(\cdot, \cdot)$ is defined only for feasible partitions $\Pi \in P(G)$. We let $V(N)$ denote $V(N, N)$.

Extending the notion of the core to PFGs is not straightforward. This is because, in the presence of externalities, a potential set of deviators S needs to make an assumption on how the rest of the agents behave (i.e., what is the partition that will form in $N \setminus S$) once they deviate. In this section, we focus on the two most common approaches found in the literature, a pessimistic and an optimistic one (see [16] for an overview of PFG solution concepts).

3.1 The Pessimistic Core

We start by defining the core in the case that deviators have a pessimistic view regarding the reaction of the remaining players. For this we need a notion of dominance.

For a given coalition structure $\Pi = (S_1, \dots, S_k) \in P(G)$, a payoff allocation $\mathbf{x} = (x_1, \dots, x_n)$ is *feasible* for Π if :

$$\sum_{j \in S_i} x_j \leq V(S_i, \Pi), \quad i = 1, \dots, k$$

Let $\Phi(\Pi)$ denote the set of feasible payoff vectors of a partition Π and let $\Phi = \bigcup_{\Pi \in P(G)} \Phi(\Pi)$. We also define the pessimistic value of a set $S \in \mathcal{F}(G)$ to be $\hat{v}(S) = \min_{\Pi \in P(G), S \in \Pi} V(S, \Pi)$.

Consider two payoff vectors \mathbf{x} and \mathbf{x}' of Φ and a coalition S . We say that \mathbf{x}' *dominates \mathbf{x} via S* if (i) $\sum_{j \in S} x'_j \leq \hat{v}(S)$ and (ii) $x'_j > x_j$ for all $j \in S$. The idea behind this type of dominance is that the agents of S are not content with \mathbf{x} because there exists another allocation (possibly on a different partition), by which they all get better off without exceeding the total payoff that is guaranteed to S in the worst case partition (which is $\hat{v}(S)$).

We will say that \mathbf{x}' *pessimistically dominates \mathbf{x}* and will write $\mathbf{x}' \text{ dom}_{pes} \mathbf{x}$ simply if there exists S such that \mathbf{x}' dominates \mathbf{x} via S . We define the *pessimistic core (p-core)* to be the set of vectors of Φ that are not (pessimistically) dominated by any other member of Φ . This is an extension of the pessimistic core of [8], which was

defined for domains without structure and with the grand coalition as the social welfare-maximizing partition.

$$p\text{-core} = \{(\Pi, \mathbf{x}) : \Pi \in P(G), \mathbf{x} \in \Phi(\Pi), \nexists \mathbf{x}' \in \Phi \text{ s.t. } \mathbf{x}' \text{ dom}_{pes} \mathbf{x}\}$$

Given a PFG on a graph G , we will associate with it the following CFG: the set of players and the graph is the same and the characteristic function is $\hat{v}(S)$. The core of this CFG is denoted by:

$$C(\hat{v}, G) = \{(\Pi, \mathbf{x}) : \Pi \in P(G), x(S) \geq \hat{v}(S) \forall S \in \mathcal{F}(G) \wedge \mathbf{x} \in I(\Pi)\}$$

where $I(\Pi)$ denotes the set of imputations in Π . We now provide a relationship between p -core and $C(\hat{v}, G)$.

THEOREM 11. *For any PFG on a graph G , let (N, \hat{v}, G) be the corresponding (pessimistic) CFG. Then $C(\hat{v}, G) \subseteq p\text{-core}$.*

PROOF. Consider a tuple (Π, \mathbf{x}) that belongs to $C(\hat{v}, G)$. Suppose that $(\Pi, \mathbf{x}) \notin p\text{-core}$. Then either \mathbf{x} is not feasible for Π or \mathbf{x} is dominated by some other feasible vector \mathbf{y} . But since $(\Pi, \mathbf{x}) \in C(\hat{v})$, then $\mathbf{x} \in \Phi(\Pi)$, therefore the only possibility is that \mathbf{x} is dominated by some other feasible vector. Hence, by definition, there exists a vector $\mathbf{y} \in \Phi$ and a set $S \in \mathcal{F}(G)$ s.t. \mathbf{y} dominates \mathbf{x} via S . That is, $y_j > x_j$ for all $j \in S$ and $\sum_{j \in S} y_j \leq \hat{v}(S)$. Hence $\sum_{j \in S} x_j < \sum_{j \in S} y_j \leq \hat{v}(S)$, which is a contradiction with $(\Pi, \mathbf{x}) \in C(\hat{v}, G)$. \square

However, the reverse direction is not generally true, as the following example demonstrates:

EXAMPLE 1. *Consider 4 players placed on a line starting from 1 up to agent 4. The idea is to setup the numbers so that the socially optimal partition is $\Pi^* = \{\{12\}, \{34\}\}$, with $V(\{12\}, \Pi^*) = V(\{34\}, \Pi^*) = 10$ and $\hat{v}(\{12\}) < 10$, $\hat{v}(\{34\}) < 10$. Then one can check that (Π^*, \mathbf{x}) with $\mathbf{x} = (5, 5, 5, 5)$ belongs to the p -core but not to $C(\hat{v}, G)$. In some detail, $\mathbf{x} = (5, 5, 5, 5)$ cannot belong to $C(\hat{v}, G)$, as \mathbf{x} is not a valid imputation of partition Π^* in the CFG described by \hat{v} , because $\hat{v}(\{12\}) < 10$ and $\hat{v}(\{34\}) < 10$ (in a PFG setting, every S may have a different partition where its pessimistic value $\hat{v}(S)$ is achieved). On the other hand, \mathbf{x} is feasible for Π^* in this specific PFG game (and belongs to p -core).*

Note that Theorem 11 does not depend on the structure of the graph and holds for any PFG. This result is interesting for two reasons. First, it demonstrates the differences between domains where the maximum social welfare is achieved by the grand coalition (and by no other partition) and domains where this does not hold. Unlike [8], where they show that, in the former case, the pessimistic core coincides with the core of the CFG, here this is no longer true. Second, Theorem 11, combined with Theorem 1, allows us to establish the non-emptiness of the core in PFGs on trees.

THEOREM 12. *For PFGs where the underlying graph is a tree, the p -core is non-empty.*

For PFGs on trees, the problem of finding an element in the core is **NP**-hard and the membership problem is **co-NP**-complete, since the class of PFGs contains the class of CFGs. Theorem 12 allows us to use Alg. 1 to find an element of the p -core. But even for lines, the algorithm cannot be implemented in polynomial time, unlike CFGs. The problem arises as we need to compute the function $\hat{v}(\cdot)$ to run Alg. 1, because of the exponentially many partitions.

THEOREM 13. *For a PFG on a line and a singleton S , it is **NP**-hard to compute $\hat{v}(S)$.*

The same holds whenever $|S| = O(1)$. We omit the proof, which is based on viewing a partition in a line as what is known in combinatorics as an ‘‘integer composition’’ [24]. Algorithmic problems

on PFGs are, naturally, computationally more demanding, as the value of a coalition varies across the exponentially many partitions. It would be interesting to identify special classes of PFGs on trees or lines for which a p -core element can be computed in polynomial time. Finally, because of Theorem 5, p -core non-emptiness cannot be guaranteed for PFGs where the underlying graph is a cycle.

3.2 The Optimistic Core

We now consider the opposite approach. Although it is perhaps less intuitive to be optimistic about the reaction of non-deviators, optimism is a well established concept in game theory and economics, as the assumption can be quite natural in certain application domains. For instance, in security games, where players in computer networks attempt to fend off attackers, optimistic players may invest only in self-protection, hoping that the other players will contribute to the overall network protection.⁶ As we will see, having optimistic deviators results in higher expectations on their part, and thus the optimistic core may be empty, unlike what was established for the p -core in Theorem 12.

To begin, we need to define a *modified notion of dominance*, which we denote by dom_{opt} . Consider two payoff vectors \mathbf{x} and \mathbf{x}' of Φ and a coalition S . We say that $\mathbf{x}' \text{ dom}_{opt} \mathbf{x}$ (via S) if there exists a partition Π with $S \in \Pi$ such that (i) $\sum_{j \in S} x'_j \leq V(S, \Pi)$ and (ii) $x'_j > x_j$ for all $j \in S$. The idea here is that since members of S are optimistic, it suffices that they find *some* partition Π and an allocation, feasible for Π , in which they are all better-off.

The *optimistic core (o-core)* is then the set of tuples (Π, \mathbf{x}) such that \mathbf{x} is not (optimistically) dominated by any other member of Φ .

$$o\text{-core} = \{(\Pi, \mathbf{x}) : \Pi \in P(G), \mathbf{x} \in \Phi(\Pi), \nexists \mathbf{x}' \in \Phi \text{ s.t. } \mathbf{x}' \text{ dom}_{opt} \mathbf{x}\}$$

The following fact is easy to verify:

FACT 2. *$o\text{-core} \subseteq p\text{-core}$.*

We will see shortly that the reverse is not true. First we will identify a necessary condition for the o -core to be non-empty. This will be done by establishing a connection of the o -core with the following CFG: the set of players and the graph structure is the same, and, in analogy to Section 3.1, we define now the optimistic value of a set $S \subseteq N$ to be $v^*(S) = \max_{\Pi \in P(G), S \in \Pi} V(S, \Pi)$. The resulting CFG is (N, v^*, G) , and its core is denoted by:

$$C(v^*, G) = \{(\Pi, \mathbf{x}) : \Pi \in P(G), x(S) \geq v^*(S) \forall S \in \mathcal{F}(G) \wedge \mathbf{x} \in I(\Pi)\}$$

Then, for a tuple (Π, \mathbf{x}) to belong to the o -core, a necessary condition is that for every S , $\sum_{i \in S} x_i \geq v^*(S)$, as implied by the following theorem.

THEOREM 14. *For any PFG on a graph G , let (N, v^*, G) be the corresponding (optimistic) CFG. Then $o\text{-core} \subseteq C(v^*, G)$.*

PROOF. Consider a tuple (Π, \mathbf{x}) that belongs to the o -core. Suppose that $(\Pi, \mathbf{x}) \notin C(v^*, G)$. Then by definition, either $\sum_{i \in N} x_i \neq \sum_{S_j \in \Pi} v^*(S_j)$ or there exists a set S such that $\sum_{i \in S} x_i < v^*(S)$. In the first case, since $\mathbf{x} \in \Phi(\Pi)$, it follows that $\sum_{i \in N} x_i < \sum_{S_j \in \Pi} v^*(S_j)$, which implies that there exists some set $S_j \in \Pi$ for which $\sum_{i \in S_j} x_i < v^*(S_j)$. Hence we can conclude that in both cases there is some set S for which $\sum_{i \in S} x_i < v^*(S)$. But then we can construct the following vector \mathbf{y} : Let Π' be the partition where $v^*(S)$ is achieved. We can have a payoff allocation to S so that $\sum_{j \in S} y_j \leq v^*(S)$ and $y_j > x_j \forall i \in S$. Furthermore, for the remaining players, we can simply allocate payoff to them so as to ensure that for every $S' \in \Pi'$, other than S ,

⁶Harikrishna *et al.* [12] have recently introduced a cooperative model against security attacks, but without explicitly examining how the network structure influences agent cooperation decisions.

$\sum_{j \in S'} y_j \leq V(S', \Pi')$. This results in a vector $\mathbf{y} \in \Phi(\Pi')$ (and hence $\mathbf{y} \in \Phi$) such that $\mathbf{y} \text{ dom}_{opt} \mathbf{x}$ via S , which is a contradiction with the fact that $(\Pi, \mathbf{x}) \in \text{o-core}$. \square

However, the reverse direction is not generally true, as demonstrated below for games defined on lines.

THEOREM 15. *There exist PFGs, where the underlying graph is a line, for which the o-core is empty.*

PROOF. We consider a graph that forms a line from agent 1 to agent 4 (in a similar manner one can generalize the example to lines with a larger number of agents). This gives rise to 8 feasible partitions, with corresponding coalitional values as follows.

Partition	Value	Value of coalitions
$\{1\}, \{2\}, \{3\}, \{4\}$	12	(3, 3, 3, 3)
$\{1\}, \{2\}, \{3, 4\}$	11	(2, 2, 7)
$\{1\}, \{2, 3\}, \{4\}$	11	(2, 7, 2)
$\{1\}, \{2, 3, 4\}$	11	(2, 9)
$\{1, 2\}, \{3\}, \{4\}$	19	(11, 4, 4)
$\{1, 2\}, \{3, 4\}$	20	(10, 10)
$\{1, 2, 3\}, \{4\}$	18	(15, 3)
$\{1, 2, 3, 4\}$	19	19

We can show that for any of these feasible partitions, there can be no vector \mathbf{x} with $\mathbf{x} \in \Phi(\Pi)$ such that $(\Pi, \mathbf{x}) \in \text{o-core}$. To see this, consider for example the welfare maximizing partition $\Pi = (\{1, 2\}, \{3, 4\})$. For any $\mathbf{x} \in \Phi(\Pi)$, we know that $x_1 + x_2 \leq 10$. But then we can construct a vector \mathbf{y} , feasible for the partition $(\{1, 2\}, \{3\}, \{4\})$, in which $y_1 + y_2 \leq 11$ and $y_1 > x_1, y_2 > x_2$. This means that $\mathbf{y} \text{ dom}_{opt} \mathbf{x}$ and (Π, \mathbf{x}) cannot belong to the o-core. If we consider the partition $\Pi' = (\{1, 2\}, \{3\}, \{4\})$, then we know that for any $\mathbf{x} \in \Phi(\Pi')$, $x_3 + x_4 \leq 8$. But then we can construct in a similar manner a vector \mathbf{y} that is feasible for Π and dominates \mathbf{x} via the set $\{3, 4\}$. By using similar arguments, we can also conclude that for the remaining 6 partitions, we cannot find a vector that would be feasible for them and undominated. \square

On the contrary, $C(v^*, G)$ is non-empty on trees by Theorem 1. This shows that unlike the p-core, which is guaranteed to exist in PFGs over trees, the o-core is a stricter concept and cannot exist for all such games. Intuitively, this is because optimistic expectations lead to a larger set of potential deviations and eventually to the absence of stable configurations.

4. BAYESIAN PARTITION FUNCTION GAMES

In this section, we initiate the study of a model that is even more general than the usual partition function games. As mentioned, the PFG literature has mostly focused on scenarios where deviators are assumed to be either pessimistic or optimistic with respect to others' behaviour. Even in approaches that attempt to move away from these extremes, definitions of dominance rely on some degree of optimism or pessimism—as, e.g., in the *recursive core* model of [16]. In many realistic scenarios, however, information regarding the behaviour of the residual agents, when a set S decides to work on its own, might be available to the deviators in the form of a probability distribution. Such a distribution could be derived from market data available, observation of historical evidence and trends, domain knowledge and so on. Arguably, it is far more useful to multiagent systems research and practice to assume that agent uncertainty is described by a probability distribution, rather than restrict attention to just one or two possible scenarios. Apart from better reflecting real life situations, such an extension could potentially allow for Bayesian inference and learning in PFGs; something that has not been discussed in the PFG literature.

We thus proceed to define *Bayesian partition function games* as normal PFGs equipped with probability distributions specifying the likelihood of a partition emerging when a specific coalition forms.

DEFINITION 3. *A Bayesian partition function game (BPFPG) is a tuple $\mathcal{B} = \langle \mathcal{P}, Pr_S(\cdot) \rangle$, where \mathcal{P} is a PFG and for every S , $Pr_S(\cdot)$ is a probability mass function, specifying the probability $Pr_S(\Pi)$ that $\Pi \in P(G)$ emerges given that $S \in \Pi$ forms.*

For each S , it should hold that $\sum_{\Pi \ni S} Pr_S(\Pi) = 1$. Also for the grand coalition, it holds that $Pr_N(\{N\}) = 1$. For a set S , the probabilities $Pr_S(\cdot)$ reflect the beliefs of S on the reaction of the residual agents, when S is considering to deviate. Then, the *expected value* that a coalition S would receive in a potential deviation in a BPFPG is: $\tilde{v}(S) = E_\Pi[V(S, \Pi)] = \sum_{\Pi \ni S} Pr_S(\Pi)V(S, \Pi)$.

It is now reasonably straightforward to define the Bayesian PFG-core as the set of *partition-allocation pairs* with efficient allocations that are weakly preferable in expectation to any potential deviation by some coalition S .

DEFINITION 4 (BPFPG-CORE). *The BPFPG-core is the set of (Π, \mathbf{x}) pairs where $\sum_{i \in C} x_i = V(C, \Pi), \forall C \in \Pi$ and for any feasible coalition, $S \in \mathcal{F}(G)$, it holds that $x(S) \geq \tilde{v}(S)$.*

We note here that, interestingly, even if (Π, \mathbf{x}) belongs to the BPFPG-core, Π may not necessarily be a welfare maximizing partition. This is because coalitions only judge whether the payoff they receive is good in expectation—i.e., at least $\tilde{v}(S)$. If an optimal partition exists that is better than Π but occurs only with a small probability, then this does not necessarily prevent the stability of Π . Though this is a departure from the usual models where the core is defined, it is very natural in a Bayesian setting.

Finding necessary and sufficient conditions for characterizing the elements of the BPFPG-core is naturally of interest. Here, we provide one *sufficient* condition for a (Π, \mathbf{x}) element to be in the BPFPG-core. First, consider the maximum attainable value $v^*(S)$ of a feasible coalition S (i.e., its best possible value under any potential partition containing S). Then, the following fact holds:

FACT 3. $v^*(S) \geq \tilde{v}(S) \forall S \in \mathcal{F}(G)$.

By Fact 3, we have:

THEOREM 16. *Let (Π, \mathbf{x}) be a Bayesian partition function game outcome with $\sum_{i \in C} x_i = V(C, \Pi), \forall C \in \Pi$. (Π, \mathbf{x}) is in the BPFPG-core if the following condition holds:*

$$x(S) \geq v^*(S), \forall S \in \mathcal{F}(G)$$

Note that this is a sufficient condition for an element to be in the BPFPG-core, without the need to take into account probability distribution estimates, even though the setting is probabilistic. Hence, if there is information available about the maximum values $v^*(S)$, then one may not need to go through computing all expectations.

This however is not a *necessary* condition: it may occur that an element (Π, \mathbf{x}) is in the core and there is an S with $v^*(S) > \sum_{i \in S} x_i$. This is demonstrated in the following example.

EXAMPLE 2. *Consider $N = \{1, 2, 3, 4\}$, and suppose $V(N) = 120$ and $\mathbf{x} = \{30, 30, 30, 30\}$. We will argue about the tuple (N, \mathbf{x}) . Consider the coalition $S = \{1, 2\}$ of agents considering to deviate, with beliefs specifying that for the two possible partitions to emerge if S breaks away, say Π_1, Π_2 , $Pr_S(\Pi_1) = Pr_S(\Pi_2) = 0.5$. Let $V(S, \Pi_1) = 100$, $V(S, \Pi_2) = 10$. Let also the value of any other feasible coalition in this game be zero in any partition. Then, since $\tilde{v}(S) = \{1, 2\} = 55$ while $x(S) = 60$, and given that all non-mentioned coalitions have zero value, it holds that for any feasible C in this setting, $x(C) \geq \tilde{v}(C)$, therefore (N, \mathbf{x}) is in the BPFPG-core. However, for the given S , $v^*(S) = 100 > x(S)$.*

In analogy to the definition of the core in superadditive CFGs, we can also define here the core with respect to the grand coalition (*BPFG-core-grand*), as the set of efficient allocations dividing up the value of N that make the deviation of a set of agents unprofitable in expectation.

DEFINITION 5 (BPFG-CORE-GRAND). *The BPFG-core-grand is the set \mathbf{x} of allocations such that $x(N) = V(N)$, and $\forall S \in \mathcal{F}(G)$, $x(S) \geq \tilde{v}(S)$.*

Since $\tilde{v}(N) = V(N)$, it is easy to verify from Def. 2 and Def. 5 that the BPFG-core-grand and the core of the CFG with characteristic function $\tilde{v}(\cdot)$ coincide.

FACT 4. *Let \mathcal{B} be a BPFG. Consider the CFG defined by the function $\tilde{v}(\cdot)$. Then, $\text{BPFG-core-grand}(\mathcal{B}) = C(\tilde{v}, G)$.*

Using Fact 4 and Theorem 1 we can now establish:

THEOREM 17. *If \mathcal{B} is a BPFG defined on a tree G , then its BPFG-core-grand is non-empty.*

The discussion right after Theorem 12, including Theorem 13, applies for $\tilde{v}(\cdot)$ as well. Hence, even though one could use Algorithm 1 to find an element of BPFG-core-grand, this cannot be done in polynomial time. It would be interesting to identify special cases of BPFGs that admit polynomial time algorithms.

We believe that BPFGs are a natural setting that deserves further exploration. Clearly, it would be interesting to obtain a correspondence between the CFG core with coalition structures and the general BPFG-core. However Theorem 17 does not hold for the general BPFG-core because the analog of Fact 4 is not always true. Namely, an element $(\Pi, \mathbf{x}) \in C(\tilde{v}, G)$ satisfies $x(C) = \tilde{v}(C)$ for every $C \in \Pi$. But this may not be a valid allocation for the BPFG since it may not hold that $x(C) = V(C, \Pi)$ (i.e., the feasibility of an allocation \mathbf{x} would have to be assessed w.r.t. Π). Hence the task of mapping the CFG core to the BPFG-core is considerably more challenging when coalition structures are involved.

5. CONCLUSIONS AND FUTURE WORK

In this paper we studied core-stability in several models of cooperative games defined on graphs that constrain the formation of coalitions. First, we obtained complexity results in the usual CFG setting, several of which are *positive* for certain graph structures of interest, such as trees and cycles which form the backbone of networks found in the real world. We then initiated the study of core-stability in PFGs defined over graphs, examining it both from a pessimistic and an optimistic viewpoint. Furthermore, we proposed a Bayesian model for PFGs, which we believe is more realistic than the usual models in economics, and suits better the coalition formation paradigms of interest to multiagent systems. We took some steps towards the study of the core in this model as well.

Regarding future work, we are particularly interested in exploring the PFG and Bayesian PFG domains further, as outlined above. We also envisage linking theoretical results in these domains to real-world applications. For instance, tractable algorithms to identify ϵ -stable coalitions could be used to inform planning decisions and optimize task execution in structured multiagent settings.

Acknowledgments

We are grateful to G. Greco, E. Malizia, L. Palopoli and F. Scarcello for extremely helpful discussions. Georgios Chalkiadakis was partially supported by the EU FP7-ICT Cognitive Systems, Interaction, and Robotics under the contract #270180 (NOPTILUS).

6. REFERENCES

- [1] R. J. Aumann and J. H. Dreze. Cooperative Games with Coalition Structures. *Intern. Journal of Game Theory*, 3(4):217–237, 1974.
- [2] Y. Bachrach, R. Meir, K. Jung, and P. Kohli. Coalitional Structure Generation in Skill Games. In *AAAI-2010*.
- [3] R. I. Brafman, C. Domshlak, Y. Engel, and M. Tennenholtz. Transferable utility planning games. In *AAAI-2010*.
- [4] G. Demange. Intermediate Preferences and Stable Coalition Structures. *J. of Mathematical Economics*, 23:45–58, 1994.
- [5] G. Demange. On Group Stability in Hierarchies and Networks. *Journal of Political Economy*, 112(4):754–778, August 2004.
- [6] X. Deng and C. Papadimitriou. On the complexity of cooperative solution concepts. *Math. of Operation Research*, 19:257–266, 1994.
- [7] E. Elkind, G. Chalkiadakis, and N. R. Jennings. Coalition structures in weighted voting games. In *ECAI-2008*, pages 393 – 397, 2008.
- [8] Y. Funaki and T. Yamato. The core of an economy with a common pool resource: A partition function form approach. *International Journal of Game Theory*, 28(2):157–171, 1999.
- [9] G. Greco, E. Malizia, L. Palopoli, and F. Scarcello. On the Complexity of Compact Coalitional Games. In *IJCAI-2009*.
- [10] G. Greco, E. Malizia, L. Palopoli, and F. Scarcello. On the complexity of the core over coalition structures. In *IJCAI-2011*.
- [11] G. Greco, E. Malizia, L. Palopoli, and F. Scarcello. On the complexity of core, kernel, and bargaining set. *Artificial Intelligence*, 175(12-13):1877–1910, 2011.
- [12] N. Harikrishna, V. Venkatanathan, and C. Pandu Rangan. Towards a cooperative defense model against network security attacks. In *Workshop on the Economics of Information Security (WEIS)*, 2010.
- [13] S. Jeong and Y. Shoham. Marginal Contribution Nets: A Compact Representation Scheme for Coalitional Games. In *Proc. of ACM-EC’05*, pages 193 – 202, 2005.
- [14] M. Jackson. A Survey of Models of Network Formation: Stability and Efficiency. In G. Demange and M. Wooders, editors, *Group Formation in Economics; Networks, Clubs and Coalitions*, 2004.
- [15] S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated Equilibria in Graphical Games. In *ACM-EC’03*, pages 42 – 47, 2003.
- [16] L. Koczy. A recursive core for partition function form games. *Theory and Decision*, 63(1):41–51, August 2007.
- [17] M. Le Breton, G. Owen, and S. Weber. Strongly Balanced Cooperative Games. *Int. Journal of Game Theory*, 20:419–427, 1992.
- [18] R. Meir, J. S. Rosenschein, and E. Malizia. Subsidies, stability, and restricted cooperation in coalitional games. In *IJCAI-2011*, pages 301–306, Barcelona, 2011.
- [19] T. Michalak, D. Marciniak, M. Szamotulski, T. Rahwan, M. Wooldridge, P. McBurney, and N. R. Jennings. A logic-based representation for coalitional games with externalities. In *AAMAS-2010*, pages 125–132, 2010.
- [20] R. Myerson. Graphs and Cooperation in Games. *Mathematics of Operations Research*, 2(3):225–229, 1977.
- [21] R. Myerson. *Game Theory: Analysis of Conflict*. 1991.
- [22] T. Rahwan, T. Michalak, N. R. Jennings, M. Wooldridge, and P. McBurney. Coalition Structure Generation in Multi-Agent Systems With Positive and Negative Externalities. In *IJCAI-09*, 2009.
- [23] P. P. Shenoy. On coalition formation: A game-theoretical approach. *International Journal of Game Theory*, 8(3):133–164, 1979.
- [24] R. Stanley. *Enumerative Combinatorics (Vol. 1)*. Cambridge University Press, 1997.
- [25] R. Thrall and W. Lucas. n-Person Games in Partition Function Form. *Naval Research Logistics Quarterly*, 10(1):281–298, 1963.
- [26] R. van den Brink. On Hierarchies and Communication. Tinbergen Institute and Free University Discussion Paper 06/056-1, 2006.