On the Complexity of Compact Coalitional Games

Gianluigi Greco
Dipartimento di Matematica
Università della Calabria
I-87036 Rende, Italy
ggreco@mat.unical.it

Enrico Malizia and **Luigi Palopoli** and **Francesco Scarcello**
D.E.I.S.
Università della Calabria
I-87036 Rende, Italy
{emalizia,palopoli,scarcello}@deis.unical.it

Abstract

A significantly complete account of the complexity underlying the computation of relevant solution concepts in compact coalitional games is provided. The starting investigation point is the setting of *graph games*, about which various long-standing open problems were stated in the literature. The paper gives an answer to most of them, and in addition provides new insights on this setting, by stating a number of complexity results about some relevant generalizations and specializations. The presented results also pave the way towards precisely carving the tractability frontier characterizing computation problems on compact coalitional games.

1 Introduction

A *coalitional game* is played by a set N of players that can form *coalitions* in order to guarantee themselves some advantage. Each coalition $S \subseteq N$ is assigned a certain worth $v(S)$ that players in S can obtain if collaborating with each other, and the outcome of the game is a vector of payoffs $(x_i)_{i \in N} \in \mathbb{R}^{|N|}$ that is meant to specify the distribution of the worth granted to each player in N.

A fundamental problem for coalitional games is to characterize the most desirable outcomes in terms of appropriate notions of fair worth distributions, which are usually called *solution concepts*. Traditionally, this issue has been studied in economics and game theory in the light of providing arguments and counterarguments about why such proposals are reasonable mathematical renderings of the intuitive concepts of fairness and stability. For instance, well-known and widely-accepted solution concepts are the *Shapely value*, the *core*, the *kernel*, the *bargaining set*, and the *nucleolus*.

More recently, however, Deng and Papadimitriou [1994] re-considered the definition of such concepts from a computer science perspective, by arguing that decisions taken by realistic agents should not involve unbounded resources to support reasoning, and by suggesting to formally capture such a bounded rationality principle by assessing the amount of resources needed to compute solution concepts in terms of their computational complexity. In particular, they considered the setting of *graph games*, where worths for coalitions over a set N of players are constructed over a weighted undirected graph $G = \langle (N, E), w \rangle$, whose nodes in N correspond to the players and where the list w encodes the edge weighting function, so that $w(e) \in \mathbb{R}$ is the weight associated with the edge $e \in E$. Then, the worth of an arbitrary coalition $S \subseteq N$ is defined as the sum of the weights associated with the edges contained in S, i.e., as the value $v(S) = \sum_{e \in E | e \subseteq S} w(e)$.

Within the setting of graph games, Deng and Papadimitriou characterized the intrinsic complexity of various tasks related to solution concepts. For instance, they showed that checking whether the core is non-empty is co-NP-complete. Moreover, they provided a polynomial-time computable closed-form characterization for the Shapely value, and showed that this value coincides with the nucleolus, whenever each of its components is non-negative. However, the picture of the complexity issues arising with graph games depicted by Deng and Papadimitriou [1994] is not complete. In particular, the complexity of the problems of checking whether an outcome belongs to the bargaining set and to the kernel was not derived and, in fact, these questions were explicitly stated as open problems there. In addition, the complexity of checking whether an outcome belongs to the nucleolus for general graph games was not derived.

1.1 Contributions

As a matter of fact, the modeling perspective introduced by Deng and Papadimitriou [1994] was very influential in the subsequent literature, and it inspired complexity analysis of several other kinds of *compactly specified* games (see, e.g., [Ieong and Shoham, 2005; Conitzer and Sandholm, 2006; Elkind *et al.*, 2007; Bilbao, 2000]), that is, of games where associations of coalitions with their worths are implicitly represented as to avoid the exponential blow-up that would result if all of them were explicitly listed. Yet, despite the prominence gained in the last few years of studying the computational properties of solution concepts in graph games, and more generally for compact games, the questions raised by Deng and Papadimitriou are still open problems.

The first contribution of this paper is precisely to solve these open problems, by showing that, for graph games,

(1) Checking whether an outcome is in the kernel is Δ_2^{P}-complete;

(2) Checking whether an outcome is the nucleolus is Δ_2^{P}-complete; and,

(3) Checking whether an outcome is in the bargaining set is $\Pi_2^{\rm P}$-complete.

These main achievements come, however, not alone in the research reported below. Indeed, we go beyond and study various computational issues arising in relevant *generalizations* and *specializations* of the setting of graph games.

Generalizations. With respect to generalizing graph games, we consider the fairly more general classes of games in polynomial characteristic function form within the setting discussed in [Bilbao, 2000], where the worth function is provided by means of an oracle, whose call requires polynomial time in size of the game representation. We say that these games are compact (or in compact form). E.g., graph games are an instance of this setting where, for each coalition $S \subseteq N$, $v(S)$ can be computed in polynomial time as a summation of the weights of all those edges $e \subseteq S$. Other notable examples of classes of compact games includes the marginal contribution nets [Ieong and Shoham, 2005] and the weighted threshold games [Elkind *et al.*, 2007]. Within this setting,

(4) We show that nothing has to be paid for the succinctness of the specifications, since all the membership results that hold for graph games also hold for any class of compact games. Indeed, we show that, on compact games, checking the membership in the kernel, bargaining set and nucleolus are still in $\Delta_2^{\rm P}$, $\Pi_2^{\rm P}$, and $\Delta_2^{\rm P}$, respectively.

Note that, while **(1)**, **(2)**, and **(3)** are mainly combinatorial contributions (for they rely on the definition of rather elaborate reductions), the contribution in **(4)** appears algorithmic in its nature. Indeed, it is achieved by showing that the various solution concepts on compact games can ultimately be defined in terms of suitable linear programs over exponentially many inequalities (succinctly specified, in their turn). In particular, those results are established by providing complexity bounds on several problems related to succinctly specified linear programs, which are of interest on their own, and where the complexity for programs with polynomially many constraints is well-known instead.

Specializations. Finally, in the last part of the paper, we investigate suitable specializations of graph games, by looking for tractable classes based on exploiting some of their graph invariants. In particular, we focused on the acyclicity property, motivated by the fact that many NP-hard problems in different application areas are known to be efficiently solvable when restricted to instances whose underlying structures can be modeled via *acyclic* graphs or *nearly-acyclic* ones, such as those graphs having *bounded treewidth*. In fact, Ieong and Shoham [2005] have already observed that deciding the membership of an outcome into the core and deciding the non-emptiness of the core are feasible in polynomial time on bounded treewidth marginal contribution nets. Here, we continue along this line of research, and

(5) We show that, on graphs having bounded treewidth, the problems of checking whether an outcome is in the kernel is feasible in polynomial time.

Interestingly, the above result is established by showing how this solution concept can be expressed in terms of an optimization problem over *Monadic Second Order Logic (MSO)* formulae. This was not observed in earlier literature, neither for the kernel, nor for other solution concepts. Thus, on graphs having bounded treewidth, tractability emerges as a consequence of Courcelle's Theorem [Courcelle, 1990] and of its generalization to optimization problems due to Arnborg, Lagergren, and Seese [Arnborg *et al.*, 1991].

2 Cooperative Game Theory

A coalitional game $\mathcal{G}$ is a pair $\langle N, v\rangle$, where N is the set of all the players and where $v : 2^N \mapsto \mathbb{R}$ is the worth function.

A vector $(x_i)_{i\in N}$ (with $x_i \in \mathbb{R}$) is an *imputation* of $\mathcal{G}$ if $\sum_{i\in N} x_i = v(N)$ and $x_i \geq v(\{i\})$, for each $i \in N$. The set of all the imputations of $\mathcal{G}$ is denoted by $X(\mathcal{G})$.

Several *(solution) concepts* have been proposed in the literature to characterize the most desirable imputations of coalitional games. Below, we recall the notions of *kernel*, *bargaining set*, and *nucleolus* (see, e.g., [Osborne and Rubinstein, 1994]), which will be the subject of our research here.

For any pair of players i and j of $\mathcal{G}$, we denote by $\mathcal{I}_{i,j}$ the set of all coalitions containing player i but not player j. The *excess* of the generic coalition S at the imputation $x \in X(\mathcal{G})$, denoted by $e(S, x)$, is defined as $v(S) - x(S)$, where $x(S)$ is a shorthand for the value $\sum_{i\in S} x_i$. The *surplus* $s_{i,j}(x)$ of i against j at x is $s_{i,j}(x) = \max_{S\in\mathcal{I}_{i,j}} e(S, x)$.

Definition 2.1. The *kernel* $\mathcal{K}(\mathcal{G})$ of a game $\mathcal{G} = \langle N, v\rangle$ is the set: $\mathcal{K}(\mathcal{G}) = \{x \in X(\mathcal{G}) \mid s_{i,j}(x) > s_{j,i}(x) \Rightarrow x_j = v(\{j\}), \forall i, j \in N, i \neq j\}$. □

Let x be an imputation. We say that (y, S) is an *objection of player i against player j to x* if $S \in \mathcal{I}_{i,j}$, $y(S) = v(S)$, and $y_k > x_k$ for all $k \in S$. If (y, S) is an objection of player i against player j to x, then we say i can object against j to x through S. A *counterobjection to the objection (y, S) of i against j* is a pair (z, T) where $T \in \mathcal{I}_{j,i}$, $z(T) = v(T)$, and $z_k \geq x_k$ for all $k \in T \setminus S$ and $z_k \geq y_k$ for all $k \in T \cap S$. If (z, T) is a counterobjection to the objection (y, S) of i against j, we say that j can counterobject to (y, S) through T. If there does not exist any counterobjection to (y, S), we say that (y, S) is a *justified objection*.

Definition 2.2. The *bargaining set* $\mathcal{B}(\mathcal{G})$ of $\mathcal{G}$ is the set of all imputations to which there is no justified objection. □

For any imputation x of $\mathcal{G}$, we define the vector: $\theta(x) = (e(S_1, x), e(S_2, x), \ldots, e(S_{2^n-1}, x))$, where the various excesses of all coalitions (but the empty one) are arranged in non-increasing order. Let $\theta(x)[i]$ denote the i-th element of $\theta(x)$. For a pair of imputations x and y, we say that $\theta(x)$ is *lexicographically smaller* than $\theta(y)$, denoted by $\theta(x) \prec \theta(y)$, if there exists a positive integer q such that $\theta(x)[i] = \theta(y)[i]$ for all $i < q$ and $\theta(x)[q] < \theta(y)[q]$.

Definition 2.3. The *nucleolus* $\mathcal{N}(\mathcal{G})$ of a game $\mathcal{G}$ is the set $\mathcal{N}(\mathcal{G}) = \{x \in X(\mathcal{G}) \mid \nexists y \in X(\mathcal{G}) \text{ s.t. } \theta(y) \prec \theta(x)\}$. □

It is well-known that, for any game $\mathcal{G}$, $\mathcal{N}(\mathcal{G})$ is a singleton (see, e.g., [Osborne and Rubinstein, 1994]).

3 On the Hardness of Solution Concepts

In this section, we shall characterize the complexity of the kernel, of the nucleolus, and of the bargaining set of graph games. With a small abuse of notation, if $\mathcal{G}$ is a weighted graph, we denote by $\mathcal{G}$ its associated graph game, too.

We next show a reduction from the Δ_2^{P}-hard problem of deciding whether the lexicographically least significant variable is true in the *lexicographically maximum satisfying assignment* for a formula [Krentel, 1986] to the problem of deciding whether an imputation belongs to the kernel of a graph game.

Theorem 3.1. *Let $\mathcal{G}$ be a graph game, and x an imputation of $\mathcal{G}$. Then, deciding whether x belongs to $\mathcal{K}(\mathcal{G})$ is Δ_2^{P}-hard.*

Proof Idea. Let $\phi = c_1 \wedge \ldots \wedge c_m$ be a **3CNF** Boolean formula, that is, a Boolean formula in conjunctive normal form over the set of lexicographically ordered variables $\{\alpha_1, \ldots, \alpha_n\}$ where each clause contains three literals (positive or negated variables) at most. Based on ϕ, we build in polynomial time the weighted graph $\boldsymbol{K}(\phi) = \langle (N_{\mathbf{K}}, E_{\mathbf{K}}), w\rangle$, as discussed next (see Figure 1 for an illustration).

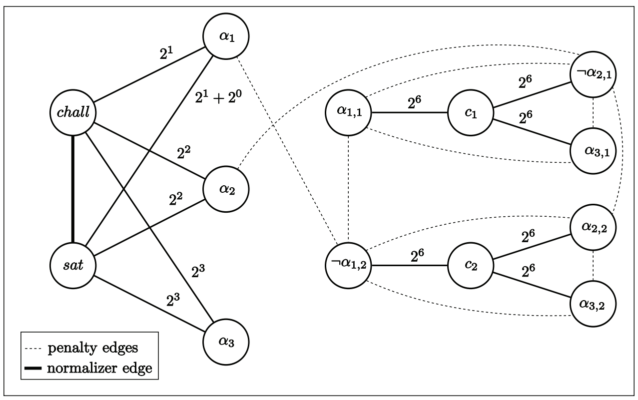

Figure 1: The game $\boldsymbol{K}(\widehat{\phi})$, where $\widehat{\phi} = (\alpha_1 \vee \neg\alpha_2 \vee \alpha_3) \wedge (\neg\alpha_1 \vee \alpha_2 \vee \alpha_3)$.

The set $N_{\mathbf{K}}$ of nodes (i.e, players) includes: a *variable player* α_i, for each variable α_i in ϕ; a *clause player* c_j, for each clause c_j in ϕ; a *literal player* $\ell_{i,j}$ (either $\ell_{i,j} = \alpha_{i,j}$ or $\ell_{i,j} = \neg\alpha_{i,j}$), for each literal ℓ_i ($\ell_i = \alpha_i$ or $\ell_i = \neg\alpha_i$, respectively) as occurring in c_j; and, two special players "*chall*" and "*sat*". The set $E_{\mathbf{K}}$ consists of the three types of edges:

Positive edges: an edge $\{c_j, \ell_{i,j}\}$ with $w(\{c_j, \ell_{i,j}\}) = 2^{n+3}$, for each literal ℓ_i occurring in c_j; an edge $\{chall, \alpha_i\}$ with $w(\{chall, \alpha_i\}) = 2^i$, for each $1 \leq i \leq n$; an edge $\{sat, \alpha_i\}$ with $w(\{sat, \alpha_i\}) = 2^i$, for each $2 \leq i \leq n$; the edge $\{sat, \alpha_1\}$ with $w(\{sat, \alpha_1\}) = 2^1 + 2^0$.

"Penalty" edges: an edge $\{\ell_{i,j}, \ell_{i',j}\}$ with $w(\{\ell_{i,j}, \ell_{i',j}\}) = -2^{m+n+7}$, for each pair of literals ℓ_i and $\ell_{i'}$ occurring in c_j; an edge $\{\alpha_{i,j}, \neg\alpha_{i,j'}\}$ with $w(\{\alpha_{i,j}, \neg\alpha_{i,j'}\}) = -2^{m+n+7}$, for each variable α_i occurring positively in c_j and negated in $c_{j'}$; an edge $\{\alpha_i, \neg\alpha_{i,j}\}$ with $w(\{\alpha_i, \neg\alpha_{i,j}\}) = -2^{m+n+7}$, for each variable α_i occurring negated in c_j.

"Normalizer" edge: $\{chall, sat\}$ with $w(\{chall, sat\}) = 1 - \sum_{e \in E_{\mathbf{K}} | e \neq \{chall, sat\}} w(e)$.

Consider the imputation x that assigns 0 to all players of $\boldsymbol{K}(\phi)$, but to sat, which receives 1. By Definition 2.1, since sat is the only player that receives in x a payoff strictly greater than her worth as a singleton coalition, and since one may note that $w(\{chall, sat\}) \geq D + 1$, where D denotes the maximum worth over all the coalitions not covering the edge $\{chall, sat\}$, we have that $x \in \mathcal{K}(\boldsymbol{K}(\phi))$ if and only if $\max_{S \in \mathcal{I}_{chall,sat}} e(S, x) \leq \max_{S \in \mathcal{I}_{sat,chall}} e(S, x)$.

Now, observe that it is never convenient for a coalition to cover a penalty edge. Thus, since the formula is satisfiable, $\max_{S \in \mathcal{I}_{chall,sat}} e(S, x)$ is precisely achieved over a coalition encoding a satisfying assignment for ϕ, thereby equating the value $CS = m \times 2^{n+3} + \max_{\sigma \models \phi} \sum_{\alpha_i | \sigma(\alpha_i) = \texttt{true}} 2^i$.

The same reasoning applies to $\max_{S \in \mathcal{I}_{sat,chall}} e(S, x)$, with the difference that for each S, $sat \in S$ implies that $x(S) = 1$ holds, and that the weight associated with the edge $\{sat, \alpha_i\}$ is 2^i, for each $2 \leq i \leq n$, while $2^1 + 2^0$ for the case where $i = 1$. Thus, this maximum coincides with the value $SC = m \times 2^{n+3} + \max_{\sigma \models \phi}(|\{\alpha_1 \mid \sigma(\alpha_1) = \texttt{true}\}| + \sum_{\alpha_i | \sigma(\alpha_i) = \texttt{true}} 2^i) - 1$.

Finally, $CS \leq SC$ holds if and only if α_1 is true in the lexicographically maximum satisfying assignment for ϕ. $\square$

Theorem 3.2. *Let $\mathcal{G}$ be a graph game, and x an imputation of $\mathcal{G}$. Then, deciding whether x belongs to $\mathcal{N}(\mathcal{G})$ is Δ_2^{P}-hard.*

Proof Idea. The reduction is again from the problem of deciding whether α_1 is true in the lexicographically maximum satisfying assignment for such a given formula ϕ. Let ϕ be a **3CNF** Boolean formula and let $(N_{\mathbf{K}}, E_{\mathbf{K}})$ be the graph built in the proof of Theorem 3.1. Let $N_k = N_{\mathbf{K}} \setminus \{sat\}$, let $\overline{N}_k$ be the set of players $\{\overline{p} \mid p \in N_k \wedge p \neq \alpha_1\}$ containing a copy of each player in N_k but α_1, and let $N_r = \{a, \overline{a}, b, \overline{b}\}$. Moreover, let $E_k = \{\{p, q\} \mid \{p, q\} \in E_{\mathbf{K}} \wedge \{p, q\} \subseteq N_k\}$.

Based on ϕ, we define the weighted graph $\boldsymbol{N}(\phi) = \langle (N_{\mathbf{N}}, E_{\mathbf{N}}), w\rangle$, as discussed next (see Figure 2 for an illustration of the construction).

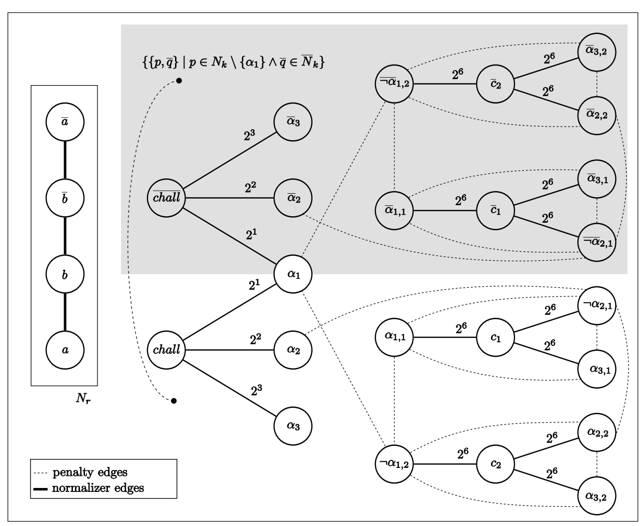

Figure 2: The game $\boldsymbol{N}(\widehat{\phi})$, where $\widehat{\phi} = (\alpha_1 \vee \neg\alpha_2 \vee \alpha_3) \wedge (\neg\alpha_1 \vee \alpha_2 \vee \alpha_3)$.

The set $N_{\mathbf{N}}$ of nodes coincides with $N_k \cup \overline{N}_k \cup N_r$. The set $E_{\mathbf{N}}$ of edges coincides with $E_k \cup \{\{\overline{p}, \overline{q}\} \mid \{p, q\} \in E_k \wedge \{\overline{p}, \overline{q}\} \subseteq \overline{N}_k\} \cup \{\{\alpha_1, \overline{q}\} \mid \{\alpha_1, q\} \in E_k \wedge \overline{q} \in \overline{N}_k\} \cup \{\{p, \overline{q}\} \mid p \in N_k \setminus \{\alpha_1\} \wedge \overline{q} \in \overline{N}_k\} \cup \{\{a, b\}, \{\overline{a}, \overline{b}\}, \{b, \overline{b}\}\}$.

Moreover, weights partition edges into three groups:

Positive edges: $w(\{c_j, \ell_{i,j}\}) = w(\{\overline{c}_j, \overline{\ell}_{i,j}\}) = 2^{n+3}$, for each literal ℓ_i occurring in c_j; $w(\{chall, \alpha_i\}) = w(\{\overline{chall}, \overline{\alpha}_i\}) = 2^i$, for each $2 \leq i \leq n$; and, $w(\{chall, \alpha_1\}) = w(\{\overline{chall}, \alpha_1\}) = 2^1$.

"Penalty" edges: $w(\{\ell_{i,j}, \ell_{i',j}\}) = w(\{\overline{\ell}_{i,j}, \overline{\ell}_{i',j}\}) = -2^{m+n+7}$, for each pair of literals ℓ_i and $\ell_{i'}$ occurring in c_j; $w(\{\alpha_{i,j}, \neg\alpha_{i,j'}\}) = w(\{\overline{\alpha}_{i,j}, \overline{\neg\alpha}_{i,j'}\}) = -2^{m+n+7}$, for each variable α_i occurring positively in c_j and negated in $c_{j'}$; $w(\{\alpha_i, \neg\alpha_{i,j}\}) = w(\{\overline{\alpha}_i, \overline{\neg\alpha}_{i,j}\}) = -2^{m+n+7}$, for each variable α_i with $i \neq 1$ occurring negated in c_j; $w(\{\alpha_1, \neg\alpha_{1,j}\}) = w(\{\alpha_1, \overline{\neg\alpha}_{1,j}\}) = -2^{m+n+7}$, for each clause c_j where α_1 negatively occurs; and $w(\{p, \overline{q}\}) = -2^{m+n+7}$, for each pair of players $p \neq \alpha_1$ and $\overline{q}$, with $p \in N_k$ and $\overline{q} \in \overline{N}_k$.

"Normalizer" edges: Consider the value $\Delta = 1 - \sum_{e \in E_{\mathbf{N}} \mid e \subseteq N_k \cup \overline{N}_k} w(e)$. Then, $w(\{a, b\}) = w(\{\overline{a}, \overline{b}\}) = \Delta + 2$ and $w(\{b, \overline{b}\}) = -\Delta - 4$.

Consider the graph game $\boldsymbol{N}(\phi)$, which can in fact be built in polynomial time, and the imputation x that assigns 0 to all players of $\boldsymbol{N}(\phi)$, but to α_1, which receives 1. Note that x is an imputation, because of the weights of the edges in N_r.

Let $S_* = \arg\max_{S \subseteq N_k} v(S)$. Along an analogous line of reasoning as in the proof of Theorem 3.1, one can show that $v(S_*) = m \times 2^{n+3} + \max_{\sigma \models \phi} \sum_{\alpha_i \mid \sigma(\alpha_i) = \mathtt{true}} 2^i$ and, in fact, that $\alpha_1 \in S_*$ if and only if α_1 evaluates to true in the lexicographically maximum satisfying assignment for ϕ.

Moreover, one can show the two following properties: (i) the excess of S_* is the maximum one, no matter what specific imputation is considered. Therefore, x is in the nucleolus if and only if it minimizes the excess of S_*. And (ii), each coalition $S \subseteq N_k$ has a *dual* coalition $\overline{S} \subseteq \overline{N}_k \cup \{\alpha_1\}$ (and viceversa) such that $\overline{S} = \{\overline{p} \mid p \in S \wedge p \neq \alpha_1\} \cup \{\alpha_1 \mid \alpha_1 \in S\}$ and $v(\overline{S}) = v(S)$.

By combining the two observations above, it emerges that if α_1 evaluates to true in the lexicographically maximum satisfying assignment for ϕ, then all the coalitions achieving the maximum excess precisely share α_1 only—in particular, this is the case because there is no single normalizer edge, which instead would have occurred in each coalition achieving the maximum excess. Hence, assigning 1 to α_1 at x is the best we can do as to minimize such maximum excesses.

On the other hand, if α_1 evaluates to false in the lexicographically maximum satisfying assignment for ϕ, then it is not convenient that α_1 retains all the worth. □

Theorem 3.3. *Let $\mathcal{G}$ be a graph game, and x be an imputation of $\mathcal{G}$. Then, deciding whether x belongs to $\mathcal{B}(\mathcal{G})$ is Π_2^{P}-hard.*

Proof Idea. Let $\Phi = (\forall\boldsymbol{\alpha})(\exists\boldsymbol{\beta})\phi(\boldsymbol{\alpha}, \boldsymbol{\beta})$ be a quantified Boolean formula over the variables $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_n\}$ and $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_r\}$, where $\phi(\boldsymbol{\alpha}, \boldsymbol{\beta}) = c_1 \wedge \ldots \wedge c_m$ is a **3CNF** formula, and where each universally quantified variable $\alpha_k \in \boldsymbol{\alpha}$ occurs only in the two clauses $c_{i(k)} = (\alpha_k \vee \neg\beta_k)$ and $c_{\bar{i}(k)} = (\neg\alpha_k \vee \beta_k)$. Deciding whether such a quantified formula is valid is known to be a Π_2^{P}-complete problem.

Based on Φ, we define the weighted graph $\boldsymbol{BS}(\Phi) = \langle (N_{\mathbf{BS}}, E_{\mathbf{BS}}), w \rangle$, as discussed next (see Figure 3).

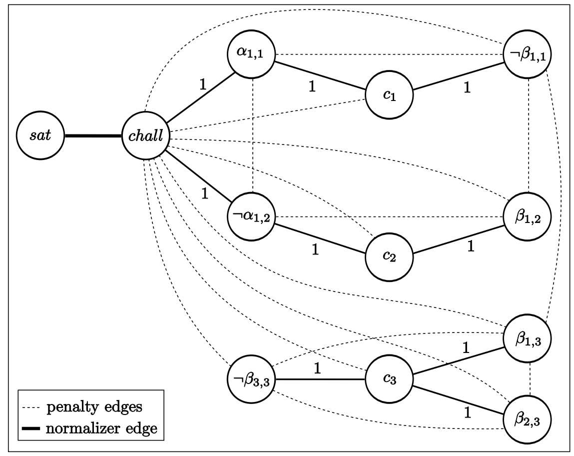

Figure 3: The game $\boldsymbol{BS}(\widehat{\Phi})$, where $\widehat{\Phi} = (\forall\alpha_1)(\exists\beta_1, \beta_2, \beta_3)(\alpha_1 \vee \neg\beta_1) \wedge (\neg\alpha_1 \vee \beta_1) \wedge (\beta_1 \vee \beta_2 \vee \neg\beta_3)$.

The set $N_{\mathbf{BS}}$ of the nodes (i.e., players) includes: a *clause player* c_j, for each clause c_j; a *literal player* $\ell_{i,j}$, for each literal ℓ_i occurring in c_j; and, two special players "*chall*" and "*sat*". The set $E_{\mathbf{BS}}$ of edges includes three kinds of edges.

Positive edges: an edge $\{c_j, \ell_{i,j}\}$ with $w(\{c_j, \ell_{i,j}\}) = 1$, for each literal ℓ_i occurring in the clause c_j; and an edge $\{chall, \ell_{i,j}\}$ with $w(\{chall, \ell_{i,j}\}) = 1$, for each literal ℓ_i of the form α_i or $\neg\alpha_i$ (i.e., built over a universally quantified variable) occurring in c_j.

"Penalty" edges: an edge $\{\gamma_{i,j}, \neg\gamma_{i,j'}\}$ with $w(\{\gamma_{i,j}, \neg\gamma_{i,j'}\}) = -m - 1$, for each variable x_i (either $\gamma_i = \alpha_i$ or $\gamma_i = \beta_i$) occurring in c_j and $c_{j'}$; an edge $\{\ell_{i,j}, \ell_{i',j}\}$ with $w(\ell_{i,j}, \ell_{i',j}) = -m - 1$, for each pair of literals ℓ_i and $\ell_{i'}$ occurring in c_j; an edge $\{chall, \ell_{i,j}\}$ with $w(\{chall, \ell_{i,j}\}) = -m - 1$, for each literal ℓ_i of the form β_i or $\neg\beta_i$ (i.e., built over an existentially quantified variable) in c_j; and, an edge $\{chall, c_j\}$ with $w(\{chall, c_j\}) = -m - 1$, for each c_j.

"Normalizer" edge: $\{chall, sat\}$ with $w(\{chall, sat\}) = n - 1 + m - \sum_{e \in E_{\mathbf{BS}} \mid e \neq \{chall, sat\}} w(e)$.

Note that the graph game $\boldsymbol{BS}(\Phi)$ can be built in polynomial time from Φ. Then, consider the imputation x that assigns m to sat, $n - 1$ to $chall$, and 0 to all other players. Note that x is indeed an imputation, since $v(N) = m + n - 1$ because $w(\{chall, sat\}) = m + n - 1 - \sum_{e \in E_{\mathbf{BS}} \mid e \neq \{chall, sat\}} w(e)$.

Firstly, we can prove that no player has a justified objection against a clause or a literal player, that no player has a justified objection against $chall$, and that no player different from $chall$ has a justified objection against sat. Therefore, we can focus on the objections of $chall$ against sat.

Then, it can be observed that the objections of $chall$ against sat are in a one-to-one correspondence with truth assignments for universally quantified variables. Moreover, counterobjections encode ways to satisfy the formula, while being consistent with the assignment encoded in the objection at hand. Therefore, x is in the bargaining set if and only if Φ is valid. $\square$

4 Membership Results

In this section, we show that the various hardness results derived in Section 3 are tight, since the corresponding membership results can be established. In fact, memberships are proven to hold over much wider classes of games than graph games, that we call *compact games*. Formally, let $\mathcal{C}$ be a class of games. We say $\mathcal{C}$ is a class of compact games if, for every game $\mathcal{G} \in \mathcal{C}$, the game encoding (whose size is denoted by $||\mathcal{G}||$) includes the set N of the players (so that, $|N| \leq ||\mathcal{G}||$), and the worth function v is implicitly given by an oracle such that, for each coalition S, $v(S)$ can be computed in polynomial time w.r.t. $||\mathcal{G}||$. Hereafter, let $\mathcal{C}_{cg}$ be a class of compact games. We start by analyzing the kernel.

Theorem 4.1. *For any $\mathcal{G} \in \mathcal{C}_{cg}$, deciding whether an imputation x belongs to $\mathcal{K}(\mathcal{G})$ is feasible in Δ_2^{P}.*

Proof. Given the imputation x, we firstly observe that, for each pair of players i and j, the value $s_{i,j}(x)$ can be computed in polynomial time by means of a binary search over the range of the possible values for the worth functions, by using an NP oracle. Indeed, for any value h in this range, we can decide in NP whether there is a coalition S such that $v(S) > h$, by guessing the coalition S, and by checking in polynomial time that $v(S) > h$. Eventually, we can compute in polynomial time the value $v(\{i\})$ for each player $i \in \{1, \ldots, n\}$. Thus, after polynomially-many oracle calls, we may check in polynomial time that for each pair of players i and j such that $x_j \neq v(\{j\})$, it is the case that $s_{i,j}(x) \leq s_{j,i}(x)$. $\square$

The cases of the bargaining set and of the nucleolus are more intricate and will be discussed in details below.

4.1 Bargaining Set

It was argued that checking whether an imputation belongs to the bargaining set is in Π_2^{P} for graph games, with the argument that one may guess in NP objections and counterobjections (hence, by focusing only on those that can be represented via polynomially many bits) [Deng and Papadimitriou, 1994]. Our main achievement is to show that membership in Π_2^{P} can indeed be established independently of the precision adopted to represent the real values of interest in the game. To this end, we first provide a useful characterization of a player i having a justified objection against some player j.

Lemma 4.2. *Player i has a justified objection against player j to x through coalition $S \in \mathcal{I}_{i,j}$ if and only if there exists a vector $y \in \mathbb{R}^{|S|}$ such that: (1) $y(S) = v(S)$; (2) $y_k > x_k$, for each $k \in S$; and, (3) $v(T) < y(T \cap S) + x(T \setminus S)$, $\forall T \in \mathcal{I}_{j,i}$.*

Theorem 4.3. *For any $\mathcal{G} \in \mathcal{C}_{cg}$, deciding whether an imputation x belongs to $\mathcal{B}(\mathcal{G})$ is feasible in Π_2^{P}.*

Proof Idea. We prove that the complementary problem of deciding whether $x \notin \mathcal{B}(\mathcal{G})$ is in Σ_2^{P}, because an NTM may guess two players i and j, and a coalition $S \in \mathcal{I}_{i,j}$, and then it may perform a co-NP check that the system of inequalities in Lemma 4.2, determined by i, j, and S, has a solution. We have to show that deciding whether this system, say $\mathtt{LP}$, has no solutions is feasible in NP. Observe that $\mathtt{LP}$ has $|S|$ variables $(y_1, \ldots, y_{|S|})$. From Helly's Theorem, given a collection $\mathcal{C} = \{c_1, \ldots, c_h\}$ of convex subsets of $\mathbb{R}^n$, $\bigcap_{c_i \in \mathcal{C}} c_i = \varnothing$ implies the existence of a collection $\mathcal{C}' \subseteq \mathcal{C}$ such that $|\mathcal{C}'| \leq n+1$ and $\bigcap_{c_i \in \mathcal{C}'} c_i = \varnothing$. Hence, if $\mathtt{LP}$ has no solutions, there is a subsystem $\mathtt{LP}'$ of $\mathtt{LP}$ consisting of (at most) $|S|+1$ inequalities that has no solutions, too. Thus, an NTM may guess $\mathtt{LP}'$, and check in polynomial time that $\mathtt{LP}'$ is infeasible, by standard linear programming techniques. $\square$

4.2 Nucleolus

Let $\mathcal{G} = \langle N, v \rangle$ be a game, and consider the following linear programming problem $\mathtt{LP}_k$, for $k > 0$:

$$\begin{aligned} \{\min \epsilon \mid\ & x(S) = v(S) - \epsilon_r \quad \forall S \in \Lambda_r, r \in \{0, \ldots, k-1\} \\ & x(S) \geq v(S) - \epsilon \quad \forall S \subseteq N, S \notin \mathcal{F}_{k-1} \\ & x \in X(\mathcal{G})\} \end{aligned}$$

where $\Lambda_0 = \{\varnothing, N\}$, $\epsilon_0 = 0$, $\mathcal{F}_0 = \varnothing$, and where ϵ_r is the optimum of $\mathtt{LP}_r$, $\Lambda_r = \{S \subseteq N \mid x(S) = v(S) - \epsilon_r,$ for every $x \in V_r\}$ with $V_r = \{x \mid x \in X(\mathcal{G}) \wedge (x, \epsilon_r)$ is an optimal solution to $\mathtt{LP}_r\}$, and $\mathcal{F}_r = \{S \subseteq N \mid x(S) = y(S), \forall x, y \in V_r\}$.

Following Maschler et al. [1979], it can be shown that there is an index $k_* \leq |N|$ such that $\mathtt{LP}_{k_*}$ has exactly one optimal solution (x_*, ϵ_{k_*}), where x_* is the nucleolus of $\mathcal{G}$.

Theorem 4.4. *For any $\mathcal{G} \in \mathcal{C}_{cg}$, computing the nucleolus of $\mathcal{G}$ is feasible in $\mathrm{F}\Delta_2^{\mathrm{P}}$, and thus deciding whether an imputation is the nucleolus of $\mathcal{G}$ is in Δ_2^{P}.*

Proof Idea. We build in $\mathrm{F}\Delta_2^{\mathrm{P}}$ a linear program $\bar{\mathtt{LP}}$ that is equivalent to $\mathtt{LP}_{k_*}$ and consists of n equalities at most, where n is the number of players N in $\mathcal{G}$. Then, the nucleolus is clearly computable in polynomial time from $\bar{\mathtt{LP}}$. To compute this program, we represent in a succinct form the above programs $\mathtt{LP}_k$, for $1 \leq k \leq n$. Roughly, equations of the form $x(S) = v(S) - \epsilon_r$ are replaced by some basis $\mathcal{B}_{k-1}$ of the affine hull of the polytope defined by $\mathtt{LP}_{k-1}$. We can show that such a basis can be computed in $\mathrm{F}\Delta_2^{\mathrm{P}}$. Moreover, given any coalition S, there is a polynomial-time oracle giving its associated inequality in $\mathtt{LP}_k$, which amounts to decide whether the vector of $\mathbb{R}^n$ encoded by S may be obtained as a linear combination of basis vectors. If this test fails, then the inequality $x(S) \geq v(S) - \epsilon$ belongs to $\mathtt{LP}_k$. Then, we compute the optimal value for this program, and continue with the next program. We can show that solving such a succinct linear program is feasible in $\mathrm{F}\Delta_2^{\mathrm{P}}$, as well as computing the dimension of its associated polytope. After at most n steps, we get a linear program whose associated polytope has dimension 0, and we stop the procedure. It can be seen that the program consisting of all the bases computed so far is equivalent to $\mathtt{LP}_{k_*}$, and of course has a unique solution (as the polytope dimension is 0), which is the nucleolus of $\mathcal{G}$. $\square$

5 Tractable Classes of Graph Games

Monadic Second Order (MSO) Logic formulae on graphs are built from logical connectives $\vee$, $\wedge$, and $\neg$, the membership relation $\in$, the quantifiers $\exists$ and $\forall$, and vertices variables and vertex sets variables—in addition, it is often convenient to use symbols like $\subseteq$, $\subset$, $\cap$, $\cup$, and $\rightarrow$ with their usual meaning, as abbreviations. Courcelle [1990] considered an extension of MSO, called MSO_2, where variables for edge sets are also allowed. The fact that an MSO_2 sentence ϕ holds over a graph G is denoted by $G \models \phi$.

An important generalization of MSO_2 formulae to *optimization* problems is presented by Arnborg et al. [1991]. Next, we state a simplified version of these kinds of problems. Optimization problems are defined over MSO_2 formulae containing free variables, and over graphs that are weighted on both nodes and edges.

Let $G = \langle (N, E), f_N, f_E \rangle$ be a weighted graph where f_N and f_E are the lists of weights associated with nodes and edges, respectively. Then, $f_N(v)$ (resp., $f_E(e)$) denotes the weight associated with $v \in N$ (resp., $e \in E$).

Let $\phi(X, Y)$ be an MSO_2 formula over the graph (N, E), where X and Y are the free variables occurring in ϕ, with X (resp., Y) being a vertex (resp., edge) set variable. For a pair of interpretations $\langle z_N, z_E \rangle$ mapping X to subsets of N and Y to subsets of E, we denote by $\phi[\langle z_N, z_E \rangle]$ the MSO_2 formula (without free variables) where X and Y are replaced by the sets $z_N(X)$ and $z_E(Y)$, respectively. A *solution* to ϕ over G is a pair of interpretations $\langle z_N, z_E \rangle$ such that $(N, E) \models \phi[\langle z_N, z_E \rangle]$ holds. The *cost* of $\langle z_N, z_E \rangle$ is the value $\sum_{x \in z_N(X)} f_N(x) + \sum_{y \in z_E(Y)} f_E(y)$. A solution of minimum cost is said *optimal*.

For a positive constant k, let hereafter $\mathcal{C}_k$ be a class of graphs having treewidth bounded by k.

Theorem 5.1 (simplified from [Arnborg *et al.*, 1991]). *Let ϕ be a fixed* MSO_2 *sentence and let $G = \langle (N, E), f_N, f_E \rangle$ be a weighted graph such that $(N, E) \in \mathcal{C}_k$. Then, computing an optimal solution to ϕ over G is feasible in* P *(w.r.t. $||G||$).*

5.1 MSO_2 and the Kernel

We are now in the position of stating a tractability result about the kernel over bounded treewidth graph games.

Theorem 5.2. *Let $\mathcal{G} = \langle (N, E), w \rangle$ be a graph game such that $(N, E) \in \mathcal{C}_k$, and let x be an imputation of $\mathcal{G}$. Then, deciding whether $x \in \mathcal{K}(\mathcal{G})$ is feasible in* P.

Proof Idea. Firstly, consider the problem of computing the coalition over which the maximum excess at x is achieved. For each $X \subseteq N$ and $Y \subseteq E$, consider the following MSO_2 formula, stating that Y is the set of all those edges $e \in E$ such that $e \subseteq X$: $proj(X, Y) \equiv \forall v, v'(\{v, v'\} \in Y \rightarrow \{v, v'\} \subseteq X) \wedge \forall v, v'(\{v, v'\} \subseteq X \wedge \{v, v'\} \in E \rightarrow \{v, v'\} \in Y)$.

Let w_E and w_N be such that $w_E(\{v, v'\}) = -w(\{v, v'\})$ and $w_N(v) = x_v$, and observe that $\max_{S \subseteq N} e(S, x) = \min_{S \subseteq N}(x(S) - v(S))$ coincides with the cost of an optimal solution to $proj(X, Y)$ over $\langle (N, E), w_N, w_E \rangle$.

Recall, now, that $x \in \mathcal{K}(\mathcal{G})$ if and only if, for each pair of players $i \neq j$, $s_{i,j}(x) > s_{j,i}(x) \Rightarrow x_j = v(\{j\})$, where $s_{i,j}(x) = \max_{S \in \mathcal{I}_{i,j}} e(S, x)$. In fact, one may modify the weights of i and j in $\mathcal{G}$ so that, in any optimal solution of the above formula, $i \in X$ and $j \notin X$, so that $\max_{S \subseteq N} e(S, x)$ coincides with $s_{i,j}$. Hence, by Theorem 5.1 and the above MSO_2 formula, $s_{i,j}$ (and $s_{j,i}$, too) is computable in polynomial time. By checking this condition for each pair i, j, membership of x in $\mathcal{K}(\mathcal{G})$ can be decided in polynomial time. □

6 Conclusions

In this paper, we have provided a picture of the complexity issues arising from graph games (by closing several long-standing open problems) and, more generally, with succinctly specified coalitional games. Our membership results apply to most of the classes of compact games proposed in the literature, while hardness results represent lower bounds (and, in fact, exact bounds) for the complexity of reasoning over them.

The paper also introduced a logic-based approach to isolate classes of tractable games for the kernel. An avenue of further research is to apply this approach to trace the tractability frontier for the nucleolus and the bargaining set.

References

[Arnborg *et al.*, 1991] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2): 308–340, 1991.

[Bilbao, 2000] J. M. Bilbao. *Cooperative Games on Combinatorial Structures*, Kluwer Academinc Publishers, 2000.

[Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence*, 170(6-7): 607–619, 2006.

[Courcelle, 1990] B. Courcelle. Graph rewriting: An algebraic and logic approach. *Handbook of TCS, Vol. B*, pp. 193–242, 1990.

[Deng and Papadimitriou, 1994] X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2): 257–266, 1994.

[Elkind *et al.*, 2007] E. Elkind, L.A. Goldberg, P.W. Goldberg, and M. Wooldridge. Computational complexity of weigthed threshold games. In *Proc. of AAAI'07*, pp. 718–723, 2007.

[Ieong and Shoham, 2005] S. Ieong and Y. Shoham. Marginal contribution nets: a compact representation scheme for coalitional games. In *Proc. of EC'05*, pp. 193–202, 2005.

[Maschler *et al.*, 1979] M. Maschler, B. Peleg, and L.S. Shapley. Geometric Properties of the Kernel, Nucleolus, and Related Solution Concepts. *Mathematics of Operations Research*, 4(4): 303–338, 1979.

[Krentel, 1986] M. W. Krentel. The complexity of optimization problems. In *Proc. of STOC'86*, pp. 69–76, 1986.

[Osborne and Rubinstein, 1994] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.