

Infeasibility Certificates and the Complexity of the Core in Coalitional Games

Enrico Malizia*

ICAR-CNR

Via P. Bucci 41/c, I-87030, Rende, Italy
malizia@si.deis.unical.it

Luigi Palopoli[†] and Francesco Scarcello[‡]

D.E.I.S., Università della Calabria

Via P. Bucci 41/c, I-87030, Rende, Italy
{palopoli, scarcello}@deis.unical.it

Abstract

This paper characterizes the complexity of the core in coalitional games. There are different proposals for representing coalitional games in a compact way, where the worths of coalitions may be computed in polynomial time. In all those frameworks, it was shown that core non-emptiness is a co-NP-hard problem. However, for the most general of them, it was left as an open problem whether it belongs to co-NP or it actually is a harder problem. We solve this open problem in a positive way; indeed, we are able to show that, for the case of transferable payoffs, the problem belongs to co-NP for any compact representation of the game where the worths of coalitions may be computed in polynomial time (also, non-deterministic polynomial time), encompassing all previous proposals of this kind. This is proved by showing that games with empty cores have small infeasibility certificates. The picture is completed by looking at coalitional games with non-transferable payoffs. We propose a compact representation based on marginal contribution nets. Also in this case, we are able to settle the precise complexity of core non-emptiness, which turns out to be Σ_2^P -complete.

1 Introduction

Coalitional games model situations where groups of players can cooperate in order to obtain certain worths, and have been extensively used to study applicative scenarios in economics and social sciences [Aumann and Hart, 1992]. Also, coalitional games are interesting in distributed AI, multi-agent systems and electronic commerce [Jeong and Shoham, 2005; Conitzer and Sandholm, 2004].

In coalitional games, a nonempty set of players joining together is called a *coalition*. The coalition including all the

players is called *grand-coalition*. The players know the worth that any coalition would get. A *feasible solution* for a coalitional game is an allowed way to assign worths (also *payoffs*) to all players. In the literature, a number of definitions of a feasible solution have been described. Each of them propose some way to assign worths to single players. Note that, in this formal context, the actions taken by players are not modeled.

There are two basic types of coalitional games [Osborne and Rubinstein, 1994]: Coalitional Games *with transferable payoffs* (or TU Games) and Coalitional Games *with non-transferable payoffs* (or NTU Games). In the former type of games, players forming a coalition can obtain a certain amount of worth they can distribute among themselves. In the latter type, a coalition guarantees a specific set of consequences that assign to its players a set of possible payoffs.

For both game types, a fundamental issue is distributing payoffs amongst participating players, which mirrors in several interesting applications [Aumann and Hart, 1992]. And, in fact, several ways of distributing utilities have been proposed, which are usually referred to as *solution concepts* (see, e.g., [Aumann and Hart, 1992] for a list of definitions). One solution concept is that of the *core*, which forces distributions that are, in a sense, “stable”, that is, no subsets of players improve their worths by leaving the grand-coalition. The core, which can be seen as an analogous of the Nash equilibrium for coalitional games [Osborne and Rubinstein, 1994], is probably the most important solution concept defined for such games (see, e.g., [Aumann, 2005]). Therefore, it is a significant issue singling out those games featuring a non-empty core, that are, games where the worth distribution can be arranged in such a way that the grand-coalition is “stable”. On the other hand, analyzing the computational complexity of solution concepts of games is an important class of problems for computer science [Papadimitriou, 2001].

In order to represent and reason about coalitional games, a way to represent the associations of coalitions with their worths is needed: doing it explicitly is unfeasible, since listing all those associations would require exponential space in the number of players. In this sense, the literature proposes a number of compact representation schemes of the worth function. For instance, Papadimitriou and Deng (1994) consider TU games where players are encoded as nodes in an arc-weighted graph, and the worth of a coalition s is computed as the sum of the weights of the arcs connecting players in s .

*Partially supported by ICAR-CNR within the research lab *Logica* of the Technology District for Logistics funded by Regione Calabria and conducted by R&D.LOG Consortium.

[†]Partially supported by ICAR-CNR within the Project *Knowledge technology and innovative services* (ICT.P04.002).

[‡]Partially supported by ICAR-CNR within the Project *High quality software* (ICT.P05.002).

A more general and more expressive representation for TU games has been recently proposed by Jeong and Shoham (2005), where the encoding is done using *marginal contribution nets*, i.e., finite sets of weighted rules, where the coalition worth is given by the sum of the weights of the logical rules triggered by its members. In both these settings, checking non-emptiness of the core was shown to be co-NP-hard, whereas in the latter—more general—setting it was left as an open problem to settle its precise complexity.

In dealing with this open question, this paper provides an answer to the rather more general question of establishing the complexity of checking core non-emptiness (in both the transferable and the non-transferable payoffs settings) for all those compact game representations satisfying the (quite weak) constraint that the associated worth function is computable in FNP (and, thus, as a special case, in polynomial time, as in the two above mentioned frameworks).

Note that our representation scheme encompasses all other compact schemes we are aware of but the one described in [Conitzer and Sandholm, 2006] for TU games, where computing the worth of a coalition is harder than FNP—the interested reader is referred to that paper for a thorough overview about compact representation schemes. Note, moreover, that for some of such compact representation schemes, the complexity of checking core non-emptiness was already established [Conitzer and Sandholm, 2006; Deng and Papadimitriou, 1994] (for instance, it is shown to be co-NP-complete in the setting given in [Deng and Papadimitriou, 1994]).

In order to prove our complexity results for TU games, we show that if the game core is empty, then there exists a small infeasibility certificate that proves it, thereby showing that for games of the quite general form we consider here, the problem of checking core non-emptiness is in co-NP (Note that this provides the answer to the problem left open in [Jeong and Shoham, 2005]). This is done by providing some results about the properties of polyhedra induced by games.

Furthermore, we consider NTU games, by defining a new compact game form, which is obtained by generalizing marginal contribution nets [Jeong and Shoham, 2005] to the non-transferable payoffs setting. For such games, we are able to show that checking core non-emptiness is Σ_2^P -hard, and also that this problem can be solved in Σ_2^P even for general games of this form, thereby settling the Σ_2^P -completeness of the problem in the non-transferable payoffs setting.

2 Preliminaries

In this section we define our formal framework of reference.

2.1 Transferable Payoffs

Games of interest in this paper are formally defined next.

Definition 2.1. A *Coalitional Game with transferable payoffs* is a pair $\langle N, v \rangle$ where

- N is the finite set of players;
- v is a function that associates with every coalition s a real number $v(s)$ (the *worth* of s) ($v: 2^N \rightarrow \mathbb{R}$).

An *outcome* for a coalitional game specifies payoffs for all players. A *solution concept* is a way to select “reasonable”

outcomes for a coalitional game: the *core* is one of the best known, as it represents a stable solution, from which players have no incentive to deviate. Let $n = |N|$. A profile \bar{x} for N is a vector of reals $(\bar{x}_1, \dots, \bar{x}_n)$, which represents a possible way to assign payoffs to players. For a coalition of players $s \subseteq N$, define $\bar{x}(s) \equiv \sum_{i \in s} \bar{x}_i$. Then, \bar{x} is said a *feasible payoff profile* if $\bar{x}(N) = v(N)$, that is, payoffs distributed among the players should be equal to the payoff available for the grand-coalition. The core is defined as follows:

Definition 2.2. The *core of a coalitional game with transferable payoffs* $\langle N, v \rangle$ is the set of all feasible payoff profiles \bar{x} such that, for all coalitions $s \subseteq N$, $\bar{x}(s) \geq v(s)$.

It immediately follows from the definition above that the core is the set of all vectors $\bar{x} \in \mathbb{R}^n$ that satisfy the following 2^n inequalities:

$$\sum_{i \in s} x_i \geq v(s), \quad \forall s \subseteq N \wedge s \neq \emptyset \quad (2.1)$$

$$\sum_{i \in N} x_i \leq v(N), \quad (2.2)$$

where the last inequality, combined with its opposite in (2.1), enforces the feasibility of computed profiles.

2.2 Non-transferable Payoffs

In coalitional NTU games each coalition is associated to a set of possible outcomes or consequences:

Definition 2.3. A *Coalitional Game without transferable payoff* is a four-tuple $\langle N, X, v, (\succsim_i)_{i \in N} \rangle$, where:

- N is a finite set of players;
- X is the set of all possible consequences;
- $v: s \rightarrow 2^X$ is a function that assigns, to any coalition $s \subseteq N$ of players, a set of consequences $v(s) \subseteq X$;
- $(\succsim_i)_{i \in N}$ is the set of all preference relations \succsim_i on X , for each player $i \in N$.

It is easy to see that coalitional games with transferable payoffs can be seen as special cases of coalitional games without transferable payoffs [Osborne and Rubinstein, 1994]. Also, the definition of the core for those latter games is an extension of that given in definition 2.2:

Definition 2.4. The *core of the coalitional game without transferable payoffs* $\langle N, X, v, (\succsim_i)_{i \in N} \rangle$ is the set of all $\bar{x} \in v(N)$ such that there is no coalition $s \subseteq N$ with a $\bar{y} \in v(s)$ such that $\bar{y} \succ_i \bar{x}$ for all $i \in s$.

3 Compact Representations

We now discuss compact representation forms, beginning with marginal contribution nets [Jeong and Shoham, 2005].

Rules in a marginal contribution net are in the form

$$\text{pattern} \rightarrow \text{value}$$

where a pattern is a conjunction that may include both positive and negative literals, with each literal denoting a player. A rule is said to *apply* to a coalition s if all the player literals occurring positively in the pattern are also in s and all the

player literals occurring negatively in the pattern do not belong to s . When more than one rule applies to a coalition, the value for that coalition is given by the contribution of all those rules, i.e., by the sum of their values. Vice versa, if no rule applies to a given coalition, then the value for that coalition is set to zero by default. For example, with rules:

$$a \wedge b \rightarrow 5, \quad b \rightarrow 2, \quad a \wedge \neg b \rightarrow 3,$$

we obtain $v(\{a\}) = 3$ (the third rule applies), $v(\{b\}) = 2$ (the second rule applies), and $v(\{a, b\}) = 5 + 2 = 7$ (both the first and the second rules apply). Using this representation scheme, games can be much more succinct than the so called *characteristic form*, where all the $2^n - 1$ values of the worth function should be explicitly listed. In any case, given such a game encoding \mathcal{G} and any coalition s , the worth $v(s)$ can be computed in linear time, that is, in $O(\|\mathcal{G}\| + \|s\|)$, which is also $O(\|\mathcal{G}\|)$, where $\|\mathcal{G}\|$ denotes the size of \mathcal{G} . As observed in [Jeong and Shoham, 2005], their representation scheme is fully expressive, in that it allows to represent any TU coalitional game, and there are games where it is exponentially more succinct than previous proposals, such as the multi-issue representation of [Conitzer and Sandholm, 2004]. For completeness, note that there are games where the size of any possible marginal nets encoding has almost the same size as the characteristic form.

Next, we are going to introduce a new and general compact representation scheme of coalitional games, a scheme where it is just required the worth function to be computable in FNP, that is, computable in polynomial time by a non-deterministic Turing transducer [Papadimitriou, 1994].

Formally, let \mathcal{C} be a class of games with transferable (resp., non-transferable) payoffs as defined by a certain given encoding scheme. Define the *worth (consequence) relation* for \mathcal{C} as the set of tuples $W_{\mathcal{C}} = \{\langle \mathcal{G}, s, w \rangle \mid \mathcal{G} \in \mathcal{C}, v_{\mathcal{G}}(s) = w\}$ (resp., $W_{\mathcal{C}} = \{\langle \mathcal{G}, s, w \rangle \mid \mathcal{G} \in \mathcal{C}, w \in v_{\mathcal{G}}(s)\}$). We say that $W_{\mathcal{C}}$ is polynomial-time computable if there is a positive integer k and a deterministic polynomial time transducer M that, given any game encoding $\mathcal{G} \in \mathcal{C}$ and a coalition s of players of \mathcal{G} , outputs a value w (resp. all consequences w) such that $\langle \mathcal{G}, s, w \rangle \in W_{\mathcal{C}}$ in at most $\|\langle \mathcal{G}, s \rangle\|^k$ steps.

We say that $W_{\mathcal{C}}$ is non-deterministically polynomial-time computable if there is a positive integer k such that $W_{\mathcal{C}}$ is k -balanced and k -decidable, as defined below. A worth (consequence) relation $W_{\mathcal{C}}$ is k -balanced if $\|w\| \leq \|\langle \mathcal{G}, s \rangle\|^k$, while it is said k -decidable if there is a non-deterministic Turing machine that decides $W_{\mathcal{C}}$ in at most $\|\langle \mathcal{G}, s, w \rangle\|^k$ time. It then follows that there is a non-deterministic Turing transducer M that may compute in $O(\|\langle \mathcal{G}, s \rangle\|^k)$ time the worth $v(s)$ (resp. some consequence in $v(s)$) of any coalition s of players of \mathcal{G} . Indeed, M guesses such a value w and a witness y of the correctness of this value (note that $W_{\mathcal{C}} \in \text{NP}$), and then verifies in *deterministic* polynomial time that $\langle \mathcal{G}, s, w \rangle \in W_{\mathcal{C}}$, possibly exploiting the witness y .

Definition 3.1. Let $\mathcal{C}(\mathcal{R})$ be the class of all games encoded according to some compact representation \mathcal{R} . We say that \mathcal{R} is a (non-deterministic) polynomial-time compact representation if the *worth relation* for $\mathcal{C}(\mathcal{R})$ is (non-deterministically) polynomial-time computable.

For instance, the extension of the marginal nets framework to games with non-transferable payoffs, presented in Section 6, is a non-deterministic polynomial-time compact representation, whereas both the above mentioned schemes of [Deng and Papadimitriou, 1994] and of [Jeong and Shoham, 2005] are polynomial time representations since, given a game \mathcal{G} encoded either as a weighted graph, or as a marginal contribution net, and given any coalition s , the worth of s in \mathcal{G} can be computed in polynomial time in the size of \mathcal{G} and s .

Our membership proofs for the core non-emptiness problem will be given in the most general setting of non-deterministic polynomial-time compact representations.

4 Separating Polyhedra

Because of (2.1) and (2.2), the core of a coalitional game with transferable payoffs and n players is a polyhedral set of \mathbb{R}^n . In this section, we prove some nice properties of polyhedral sets that will be useful to deal with such games.

4.1 Preliminaries on Polyhedral Sets

We next give some useful definitions and facts about polyhedral sets. We refer the interested reader to any book on this subject for further readings (see, e.g., [Grünbaum, 1967]).

Let $n > 0$ be any natural number. A *Polyhedral Set* (or *Polyhedron*) P of \mathbb{R}^n is the intersection of a finite set S of closed halfspaces of \mathbb{R}^n . Note that in this paper we always assume, unless otherwise stated, that $n > 0$. We denote this polyhedron by $\text{Pol}(S)$ and we denote S by $\text{Half}(P)$.

Recall that a *hyperplane* H of \mathbb{R}^n is a set of points $\{x \in \mathbb{R}^n \mid a^T x = b\}$, where $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$. The closed *half-space* H^+ is the set of points $\{x \in \mathbb{R}^n \mid a^T x \geq b\}$. We say that these points *satisfy* H^+ . We denote the points that do not satisfy this halfspace by H^- , i.e., $H^- = \mathbb{R}^n \setminus H^+ = \{x \in \mathbb{R}^n \mid a^T x < b\}$. Note that H^- is an open halfspace. We say that H *determines* H^+ and H^- . Define the *opposite* of H as the set of points $\bar{H} = \{x \in \mathbb{R}^n \mid a'^T x = b'\}$, where $a' = -1 \cdot a$ and $b' = -1 \cdot b$. Note that $\bar{H}^+ = H^- \cup H$, since it is the set of points $\{x \in \mathbb{R}^n \mid a^T x \leq b\}$.

Let P be a polyhedron and H a hyperplane. Then, H *cuts* P if both H^+ and H^- contain points of P , and we say that H *passes through* P , if there is a non-empty touching set $C = H \cap P$. Furthermore, we say that H *supports* P , or that it is a *supporting hyperplane* for P , if H does not cut P , but passes through P , i.e., it just touches P , as the only common points of H and P are those in their intersection C .

Moreover, we say that H^+ is a *supporting halfspace* for P if H is a supporting hyperplane for P and $P \subseteq H^+$. Note that $P \subseteq \text{Pol}(S)$ for any set of halfspaces $S \subseteq \text{Half}(P)$, since the latter polyhedron is obtained from the intersection of a smaller set of halfspaces than P . We say that such a polyhedron is a *supporting polyhedron* for P .

Recall that, for any set $A \subseteq \mathbb{R}^n$, its dimension $\dim(A)$ is the dimension of its affine hull. For instance, if A consists of two points, or it is a segment, its affine hull is a line and thus $\dim(A) = 1$. By definition, $\dim(\emptyset) = -1$, while single points have dimension 0.

A set $F \subseteq P$ is a *face* of P if either $F = \emptyset$, or $F = P$, or if there exists a supporting hyperplane H_F of P such that F




Figure 1: Construction of an infeasibility certificate for the core.

is their touching set, i.e., $F = H_F \cap P$. In the latter case, we say that F is a *proper face* of P . A *facet* of P is a proper face of P with the largest possible dimension, that is, $\dim(P) - 1$. The following facts are well known [Grünbaum, 1967]:

1. For any facet F of P , there is a halfspace $H^+ \in \text{Half}(P)$ such that $F = H^+ \cap P$. We say that H^+ *generates* F .
2. For any proper face F of P , there is a facet F' of P such that $F \subseteq F'$.
3. If F and F' are two proper faces of P and $F \subset F'$, then $\dim(F) < \dim(F')$.

4.2 Separating Polyhedra from a Few Supporting Halfspaces

Lemma 4.1. *Let P be a polyhedron of \mathbb{R}^n with $\dim(P) = n$, and H_F^+ a supporting halfspace of P whose touching set is F . Then, there exists a set of halfspaces $\mathcal{H}_F \subseteq \text{Half}(P)$ such that $|\mathcal{H}_F| \leq n - \dim(F)$, H_F^+ is a supporting halfspace of $\text{Pol}(\mathcal{H}_F)$, and their touching set C is such that $F \subseteq C$.*

Proof. (Rough Sketch.) The proof is by induction. *Base case:* If $\dim(F) = n - 1$ we have that the touching face $F = H_F^+ \cap P$ is a facet of P . Thus, from Fact 1, F is generated by some halfspace $H^+ \subseteq \text{Half}(P)$ such that $H^+ \cap P = F$, as for H_F . Since $\dim(F) = \dim(H) = \dim(H_F) = n - 1$, it easily follows that in fact $H = H_F$ holds. Thus, H_F^+ is trivially a supporting halfspace of H^+ , and this case is proved: just take $\mathcal{H}_F = \{H^+\}$ and note that $|\mathcal{H}_F| = 1$.

Inductive step: By the induction hypothesis, the property holds for any supporting halfspace $H_{F'}^+$ of P such that its touching face F' has a dimension $d \leq \dim(F') \leq n - 1$, for some $d > 0$. We show that it also holds for any supporting halfspace H_F^+ of P , whose touching face F has a dimension $\dim(F) = d - 1$. For space limitations, we just give the proof idea, with the help of Figure 1. Since F is not a facet, from Fact 2 there exists a facet F' of P such that $F \subset F'$. In the three-dimensional example shown in Figure 1, F is the vertex at the bottom of the diamond, and F' is some facet on its “dark side.” Let $C = H_F \cap H_{F'}$, and consider the rotation of H_F about C on the opposite direction w.r.t. $H_{F'}$ that first

touches P , say $H_{F''}$. As shown in Figure 1, the face F'' —an edge of the diamond—properly includes F and its dimension is at least $d > \dim(F)$, by fact 3. It can be shown that, given such a pair of halfspaces $H_{F'}^+$ and $H_{F''}^+$, it holds that H_F^+ is a supporting halfspace of the polyhedron $\text{Pol}(H_{F'}^+, H_{F''}^+)$, which is called *roof*. Formally, the proof proceeds by exploiting the induction hypothesis. Intuitively, consider $H_{F''}$: we want a set \mathcal{H}_F supported by H_F^+ and consisting of just halfspaces taken from $\text{Half}(P)$, and $H_{F''}^+$ does not belong to this set, because it does not generate a facet of P . However, we can see that it is a supporting halfspace for $H_{F'}^+ \cap H_{F''}^+$ —the roof, which correspond to faces having higher dimension than F'' . In the running example, they are both facets of the diamond, and hence the property immediately holds (base case). In general, the procedure may continue, encountering each time at least one facet by fact 2, and one more face with a higher dimension than the current one by fact 3. Eventually, in our example we get $\mathcal{H}_F = \{H_{F'}^+, H_{F''_1}^+, H_{F''_2}^+\}$. Moreover, recall that $\dim(F') = n - 1$ and $\dim(F'') > \dim(F) = d - 1$. Then, by the induction hypothesis, $|\mathcal{H}_F| \leq |\mathcal{H}_{F'}| + |\mathcal{H}_{F''}| = 1 + |\mathcal{H}_{F''}| \leq 1 + n - \dim(F'') \leq 1 + n - d = n - \dim(F)$. \square

5 Small Emptiness Certificates for the Core

In this section, we prove our main results on TU games.

With a little abuse of notations, since coalitions correspond to the inequalities (2.1) and hence to the associated halfspaces of \mathbb{R}^n , hereafter we use these terms interchangeably.

Definition 5.1. Let $\mathcal{G} = \langle N, v \rangle$ be a game with transferable payoffs. A coalition set $S \subseteq 2^N$ is a *certificate of emptiness* (or *infeasibility certificate*) for the core of \mathcal{G} if the intersection of $\text{Pol}(S)$ with the grand-coalition halfspace (2.2) is empty.

The definition above is motivated by the following observation. Let P be the polyhedron of \mathbb{R}^n obtained as the intersection of all halfspaces (2.1). Since S is a subset of all possible coalitions, $P \subseteq \text{Pol}(S)$. Therefore, if the intersection of $\text{Pol}(S)$ with the grand-coalition halfspace (2.2) is empty, the intersection of this halfspace with P is empty, as well.

Theorem 5.2. *Let $\mathcal{G} = \langle N, v \rangle$ be a game with transferable payoffs. If the core of \mathcal{G} is empty, there is a certificate of emptiness S for it such that $|S| \leq |N|$.*

Proof. (Sketch.) Let $n = |N|$ and P be the polyhedron of \mathbb{R}^n obtained as the intersection of all halfspaces (2.1). Since we are not considering the feasibility constraint (2.2), there is no upper-bound on the values of any variable x_i , and thus it is easy to see that $P \neq \emptyset$ and $\dim(P) = n$.

Let H_P^+ be the halfspace defined by the grand-coalition inequality (2.2). If the core of \mathcal{G} is empty, the whole set of inequalities has no solution, that is, $P \cap H_P^+ = \emptyset$.

Let \bar{H}_F^+ be the halfspace parallel to H_P^+ that first touches P , that is, the smallest relaxation of H_P^+ that intersect P . Consider the opposite H_F^+ of \bar{H}_F^+ , as shown in Figure 1, on the left. By construction, $H_P^+ \cap H_F^+ = \emptyset$, $H_F = \bar{H}_F$ is a supporting hyperplane of P , and H_F^+ is a supporting halfspace of P . Let F be the touching set of H_F with P , and let $d = \dim(F)$. In Figure 1, it is the vertex at the bottom of the diamond P . From Lemma 4.1, there is a set of halfspaces $S \subseteq \text{Half}(P)$, with $|S| \leq n - d$, such that H_P^+ is a supporting halfspace for $\text{Pol}(S)$. It follows that $H_P^+ \cap \text{Pol}(S) = \emptyset$, whence S is an infeasibility certificate for the core of \mathcal{G} . Finally, note that the largest cardinality of S is n , and corresponds to the case $\dim(F) = 0$, that is, to the case where the face F is just a vertex. Therefore the maximum cardinality of the certificate is n . In our three-dimensional example, such a certificate is $\{H_{F_1}^+, H_{F_1'}^+, H_{F_2}^+\}$, as shown in Figure 1. \square

Note that the above proof is constructive and has a nice geometrical interpretation. However, for the sake of completeness, we point out that—as we have recently found—the above result on infeasibility certificates may be also obtained as a consequence of Helly’s Theorem on the intersection of families of convex sets [Danzer *et al.*, 1963], whose proof is rather different, as it relies on algebraic techniques.

Exploiting the above property, we can now state our general result on the complexity of core non-emptiness for any (non-deterministic) polynomial-time compact representation.

Theorem 5.3. *Let \mathcal{R} be a non-deterministic polynomial-time compact representation. Given any coalitional game with transferable payoffs $\mathcal{G} \in \mathcal{C}(\mathcal{R})$, deciding whether the core of \mathcal{G} is not empty is in co-NP.*

Proof. Let $\mathcal{G} = \langle N, v \rangle$ be a game with transferable payoffs. If its core is empty, from Theorem 5.2, there is an infeasibility certificate S , with $|S| \leq n$, where n is the number of players of \mathcal{G} . For the sake of presentation, let us briefly sketch the case of a polynomial-time deterministic representation \mathcal{R} . In this case, a non-deterministic Turing machine may check in polynomial time that the core is empty by performing the following operations: (i) guessing the set S , i.e., the coalitions of players corresponding to the halfspaces in S ; (ii) computing (in deterministic polynomial time) the worth $v(s)$, for each $s \in S$, and for the grand-coalition N ; and (iii) checking that $\text{Pol}(S) \cap H_P^+ = \emptyset$, where H_P^+ is the halfspace defined by the grand-coalition inequality (2.2). Note that the last step is feasible in polynomial time, as we have to solve a linear system consisting of just $n + 1$ inequalities.

The case of a non-deterministic polynomial-time compact representation \mathcal{R} is a simple variation where, at step (ii), for each $s \in S$, the machine should also guess the value $w = v(s)$ and a witness y that $\langle \mathcal{G}, s, w \rangle \in W_{\mathcal{C}(\mathcal{R})}$. \square

The above result settles the precise complexity of the core non-emptiness problem for marginal contribution nets, as asked for in [Jeong and Shoham, 2005].

Corollary 5.4. *Given a coalitional game with transferable payoffs encoded as a marginal contribution net, deciding whether its core is not empty is co-NP-complete.*

6 Non-Transferable Payoffs Increase the Complexity

For games where the payoffs cannot be transferred among the players, core non-emptiness turns out to be harder than in the case of transferable payoffs we have studied above.

First, for NTU games we next define a notation similar to that of marginal contribution nets, to describe the consequences of these games in a compact form.

Also here, games are described by associating player patterns with consequences, a coalition worth being thus characterized by the sums of the contributions of the rules whose patterns are satisfied by the players in the coalition.

Definition 6.1. *A marginal contribution net for game with non-transferable payoffs is a finite set of rules of the form*

$$\text{pattern} \rightarrow \text{consequences},$$

where *pattern* is a conjunction of positive and negative player literals, and *consequences* is a set of possible payoff addenda for the players in the coalition that trigger this rule, that is, for players occurring as positive literals in *pattern*.

Formally, given such a rule r , we say that a coalition s of players *meets* its pattern, if each player occurring positively in *pattern* also occurs in s , and none of the players occurring negatively in *pattern* occurs in s . The consequences of r are a set of vectors assigning an increment (either positive or negative) to some players occurring positively in *pattern*. All other players get no increment out of this vector. Syntactically, we thus specify only the contributions for the players to be incremented (see example below). For each player p , there is a default implicitly specified rule, which is triggered by the player p and assigns to it the (initial) value 0. Then, the only consequence of coalitions that do not meet any non-default rule is the outcome assigning payoff 0 to all their players.

Let s be a coalition and R the set of rules that s meets. The set of consequences $v(s)$ of s is the set of all imputation vectors that can be obtained by taking the sum of any tuple of vectors $\bar{x}_1, \dots, \bar{x}_{|R|}$, with each \bar{x}_i belonging to the consequences of some rule $r_i \in R$.

Example 6.2. Let us consider a game involving players a, b and c . Then, consider the rules

$$a \wedge b \rightarrow [a+=1], [a+=2, b+=1]$$

$$b \wedge \neg c \rightarrow [b+=4].$$

Then the set of consequences $v(\{a\})$ is the singleton $\{(0, 0, 0)\}$, since only the implicit default rules apply. On the other hand, $v(\{a, b\}) = \{(1, 4, 0), (2, 5, 0)\}$, $v\{b, c\} = \{(0, 0, 0)\}$, and $v(\{a, b, c\}) = \{(1, 0, 0), (2, 1, 0)\}$.

Theorem 6.3. *Let \mathcal{R} be a non-deterministic polynomial-time compact representation. Given any coalitional game with non-transferable payoffs $\mathcal{G} \in \mathcal{C}(\mathcal{R})$, deciding whether the core of \mathcal{G} is not empty is in Σ_2^P . In particular, it is Σ_2^P -complete if \mathcal{R} is the marginal contribution nets framework. However, if \mathcal{R} is a deterministic polynomial-time compact representation, the problem is in co-NP.*

Proof. (Rough Sketch.) (*Membership in Σ_2^P*). Let $\mathcal{G} = \langle N, X, v, (\succsim_i)_{i \in N} \rangle$ be a game in $\mathcal{C}(\mathcal{R})$ with non-transferable payoffs. A non-deterministic Turing machine with an oracle in NP may decide in polynomial time that the core is not empty as follows: (i) guessing the profile $w \in v(N)$ and of a witness y that w is a consequence of $v(N)$; (ii) exploiting y , checking in polynomial time that $\langle \mathcal{G}, N, w \rangle \in W_{\mathcal{C}(\mathcal{R})}$; and (iii) exploiting the oracle, checking that w belongs to the core. Indeed, it is easy to see that the latter problem is in co-NP.

(*Membership in co-NP*). If \mathcal{R} is a deterministic polynomial-time compact representation, a non-deterministic Turing machine M may decide in polynomial time that the core is empty as follows. It computes in polynomial time all consequences of the grand-coalition and guesses, for each $w \in v(N)$, a witness y_w that w is not in the core. Then, exploiting these witnesses, M checks in polynomial time that all such profiles do not belong to the core.

(*Σ_2^P -hardness*). The reduction is from the problem of deciding the validity of 2QBF formulae in 3DNF. Let $\Phi = \exists \alpha \forall \beta \phi(\alpha, \beta)$, where α and β are vectors of boolean variables, and $\phi(\alpha, \beta)$ is a boolean formula in 3DNF. From Φ , we build in polynomial time the following NTU game \mathcal{G}_Φ , encoded as a marginal contribution net. The players of \mathcal{G}_Φ are: two players a_i^T and a_i^F , called existential players, corresponding to each variable α_i of Φ ; two players b_i^T and b_i^F , called universal players, for each variable β_i of Φ ; a player d_i for each disjunct δ_i of Φ ; and two more players *sat* and *good*, where the former is related to the satisfiability of $\phi(\alpha, \beta)$, while the latter is a player that receives a penalty $-B$ in coalitions that do not behave as we would like to, where B is a fixed number larger than any value that any coalitions may achieve.

For any literal ℓ occurring in some disjunct of the formula, denote by $p(\ell)$ the corresponding player of the game. E.g., if $\ell = \neg \alpha_i$, then $p(\ell) = a_i^F$; if $\ell = \alpha_i$, then $p(\ell) = a_i^T$. Then, the consequences for coalitions in \mathcal{G}_Φ are defined through the following rules:

$$\begin{aligned} b_i^T &\rightarrow [b_i^T += 1]; & b_i^F &\rightarrow [b_i^F += 1] \\ a_i^T \wedge a_i^F &\rightarrow [a_i^T += 1], [a_i^F += 1] \\ d_i &\rightarrow [d_i += 1]; & \text{sat} &\rightarrow [\text{sat} += 1]; & \text{good} &\rightarrow [\text{good} += 1] \\ &\text{for each disjunct } \delta_i = \ell_1 \wedge \ell_2 \wedge \ell_3 \text{ of } \phi(\alpha, \beta), \\ p(\ell_1) \wedge p(\ell_2) \wedge p(\ell_3) \wedge \neg d_i \wedge \text{good} &\rightarrow [\text{good} += -B] \\ d_i \wedge \neg \text{sat} \wedge \text{good} &\rightarrow [\text{good} += -B] \\ \neg d_1 \wedge \dots \wedge \neg d_m \wedge \text{sat} \wedge \text{good} &\rightarrow [\text{good} += -B] \\ b_i^T \wedge \neg \text{sat} \wedge \text{good} &\rightarrow [b_i^T += 1, \text{good} += 1] \\ b_i^F \wedge \neg \text{sat} \wedge \text{good} &\rightarrow [b_i^F += 1, \text{good} += 1] \\ a_i^T \wedge \neg a_i^F &\rightarrow [a_i^T += 1]; & a_i^F \wedge \neg a_i^T &\rightarrow [a_i^F += 1] \end{aligned}$$

$$\begin{aligned} \neg a_i^T \wedge \neg a_i^F \wedge \text{good} &\rightarrow [\text{good} += -B] \\ \neg b_i^T \wedge \neg b_i^F \wedge \text{good} &\rightarrow [\text{good} += -B] \\ &\text{for all players } p \neq \text{good}, \quad p \wedge \neg \text{good} \rightarrow [p += -B]. \end{aligned}$$

Then, it can be shown that Φ is valid if and only if \mathcal{G}_Φ has a non empty core. First note that any imputation vector for the grand-coalition N assigns 1 to universal players, disjunct players, and to *sat* and *good*. For any pair of existential players a_i^T and a_i^F , they get either 0 or 1 but never the same value. Then, we can associate with such a configuration of existential players a truth-value assignment for the corresponding boolean variables: if \bar{x} is an imputation vector, define $\sigma_{\bar{x}}$ such that $\sigma_{\bar{x}}(\alpha_i) = \text{true}$ if a_i^T takes 0 in \bar{x} , and $\sigma_{\bar{x}}(\alpha_i) = \text{false}$ if a_i^F takes 0 in \bar{x} . Intuitively, if the formula is valid then $\sigma_{\bar{x}}$ is a witness of validity, and there is such an imputation vector \bar{x} in the core. Indeed, in this case, to improve the payoff for all of its members, a coalition should avoid to include *sat*. However, this is impossible if Φ is valid. \square

References

- [Aumann and Hart, 1992] Robert J. Aumann and Sergiu Hart, editors. *Handbook of Game Theory. Vol 1 (1992), Vol 2 (1994) and Vol 3 (2002)*. North-Holland.
- [Aumann, 2005] Robert J. Aumann. War and peace. The Nobel Foundation, 2005. Nobel Prize Lecture.
- [Conitzer and Sandholm, 2004] Vincent Conitzer and Thomas Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. *AAAI-04*, 219–225, AAAI-Press.
- [Conitzer and Sandholm, 2006] Vincent Conitzer and Thomas Sandholm. Complexity of constructing solutions in the core based on synergies among coalitions. *AIJ*, 170(6-7):607–619, 2006.
- [Danzer *et al.*, 1963] L. Danzer, B. Grünbaum, and V. Klee. Helly’s theorem and its relatives. In *Proc. of Symposia in Pure Mathematics 7*, 101–180, AMS, 1963.
- [Deng and Papadimitriou, 1994] Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative game solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.
- [Grünbaum, 1967] Branko Grünbaum. *Convex Polytopes*, volume XVI of *Pure and Applied Mathematics*. John Wiley & Sons, NY, 1967.
- [Jeong and Shoham, 2005] Samuel Jeong and Yoav Shoham. Marginal contribution nets: a compact representation scheme for coalitional games. In *Proc. of EC ’05*, 193–202, Vancouver, Canada, ACM Press, 2005.
- [Osborne and Rubinstein, 1994] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, MA, 1994.
- [Papadimitriou, 1994] Christos H. Papadimitriou. *Computational Complexity*. Addison Wesley, Reading, MA, 1994.
- [Papadimitriou, 2001] Christos H. Papadimitriou. Algorithms, games, and the Internet. *STOC*, 2001.