

Cooperative Concurrent Games

Julian Gutierrez
University of Oxford
Oxford, UK

Julian.Gutierrez@cs.ox.ac.uk

Sarit Kraus
Bar-Ilan University
Ramat Gan, Israel
sarit@cs.biu.ac.il

Michael Wooldridge
University of Oxford
Oxford, UK
mjw@cs.ox.ac.uk

ABSTRACT

In *rational verification*, one is interested in understanding which temporal logic properties will hold in a concurrent game, under the assumption that players choose strategies that form an equilibrium. Players are assumed to behave rationally in pursuit of individual goals, typically specified as temporal logic formulae. To date, rational verification has only been studied in *noncooperative* settings. In this paper, we extend the rational verification framework to *cooperative games*, in which players may form coalitions to collectively achieve their goals. We base our study on the computational model given by concurrent game structures and focus on the core as our basic solution concept. We show the core of a concurrent game can be logically characterised using ATL*, and study the computational complexity of key decision problems associated with the core, which range from problems in PSPACE to problems in 3EXPTIME. We also discuss a number of variants of the main definition of the core, leading to the issue of credible coalition formations, and a possible implementation of the main reasoning framework.

KEYWORDS

Concurrent Games; Cooperative Games; Logic; Formal Verification

ACM Reference Format:

Julian Gutierrez, Sarit Kraus, and Michael Wooldridge. 2019. Cooperative Concurrent Games. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Concurrent games have become established as a key semantic model for concurrent and multi-agent systems, in both the multi-agent systems/AI community and the verification/computer science community [1, 13, 14, 29]. A concurrent game [1] is a finite-state environment, populated by a collection of independent agents. A game operates over an infinite sequence of rounds, where at each round, each agent chooses an action to perform. Preferences in concurrent games are typically modelled by assuming that each agent is associated with a temporal logic goal formula [8], which it desires to see satisfied. The infinite plays generated by a game (modelling the computation runs of a concurrent and multi-agent system) will either satisfy or fail to satisfy each player’s goal, and since the satisfaction of a player’s goal is dependent on the choices of other players, then they must make choices strategically.

In all previous studies that we are aware of, concurrent games are assumed to be *noncooperative*: players act alone, and binding

agreements between players are assumed to be impossible. The solution concepts used in previous studies of concurrent games have therefore been noncooperative – primarily Nash equilibrium and refinements thereof. In such a noncooperative setting, the basic questions that we ask of a concurrent game are, for example, whether a particular temporal logic property holds in *some* computation of the system that could arise through players selecting strategies that form a Nash equilibrium (the E-NASH problem) or whether a property holds on *all* such computations (A-NASH) [13, 14, 29]. These problems can be understood as game-theoretic counterparts of the conventional model checking problem [7]. The complexity of decision problems surrounding such problems in concurrent games modelling the behaviour of multi-agent systems have been extensively studied, and software tools are available that support the analysis of concurrent games using these concepts [3, 15, 19].

The aim of the present paper is to extend the study of concurrent games to include *cooperative* solution concepts [20, 23]. Thus, we assume there is some (exogenous) mechanism through which players in a concurrent game can reach binding agreements and form coalitions in order to collectively achieve goals (although we emphasise that the nature of such a mechanism is beyond the scope of the present work). The possibility of binding cooperation and coalition formation eliminates some undesirable equilibria that arise in noncooperative settings, and makes available a range of outcomes that cannot be achieved without cooperation. We focus on the *core* as our key solution concept. The basic idea behind the core is that a game outcome is said to be core-stable if no subset of players could benefit by collectively deviating from it; the core of a game is the set of core-stable outcomes. Now, in conventional cooperative games (characteristic function games with transferable utility [4]), this intuition can be given a simple and natural formal definition, and as a consequence the core is probably the most widely-studied solution concept for cooperative games. However, the conventional definition of the core does not naturally map into our concurrent game setting, because in such games, coalitions are subject to *externalities*: whether or not a coalition has a beneficial deviation depends not just on the makeup of that coalition, *but on the behaviour of the remaining players* in the game too.

We begin by introducing the framework of concurrent games, and then proceed to define two variations of the core for such settings. In the first, a coalition of players are assumed to have a beneficial deviation if they have some course of action available to them which they would benefit from *no matter what the remaining players did*. This “worst case” analysis is easily defined, but requires a deviation to be beneficial against *all* courses of action by the remaining players – even those that the remaining agents would not rationally choose (*cf.*, the concept of the α -core in the game theory literature). This motivates a second definition, where a deviation is only required to be beneficial against all courses of action by

remaining players that are *credible*, in the sense that they would be no worse off than they were originally. In each case, we formally define the solution concept, identify some of its key computational properties, give logical characterisations, and where possible also complexity results, which range from properties that can be checked in PSPACE to properties that can be checked in 3EXPTIME. We also study model theoretic properties related to the core: in particular, whether it is never empty in the model of games we consider, and whether temporal logic properties hold across bisimilar systems over plays (computation runs) induced by elements in the core, a desirable property from a formal verification viewpoint.

Structure of the paper. The rest of this paper is organised as follows. In Section 2 we provide necessary background on logic and games and in Section 3 we define the core and the main computational properties associated with it. In Section 4 we present our main results and in Section 5 we study the issue of credible coalition formations, with associated complexity results. Then, in Section 6 we present some concluding remarks and related work, along with a discussion about different issues related to the core, including its implementation in practice using model checking techniques.

2 PRELIMINARIES

Given any set $S = \{s, q, r, \dots\}$, we use S^* , S^ω , and S^+ for, respectively, the sets of finite, infinite, and non-empty finite sequences of elements in S . If $w_1 = s's'' \dots s^k \in S^*$ and w_2 is any other (finite or infinite) sequence, we write $w_1 w_2$ for their concatenation $s's'' \dots s^k w_2$. For $Q \subseteq S$, we write S_{-Q} for $S \setminus Q$ and S_{-i} if $Q = \{i\}$. We extend this notation to tuples $u = (s_1, \dots, s_k, \dots, s_n)$ in $S_1 \times \dots \times S_n$, and write u_{-k} for $(s_1, \dots, s_{k-1}, s_{k+1}, \dots, s_n)$, and similarly for sets of elements, that is, by u_{-Q} we mean u without each s_k , for $k \in Q$. Given a sequence w , we write $w[t]$ for the element in position $t + 1$ in the sequence; for instance, $w[0]$ is the first element of w . We also write $w[l \dots m]$ for the sequence $w[l] \dots w[m]$, and $w[l \dots m]$ for $w[l] \dots w[m - 1]$; if $m = 0$, we let $w[l \dots m]$ be the empty sequence, denoted ϵ .

Games. Let $\text{Ag} = \{1, \dots, n\}$ be a set of *players* and St a set of *states*. For each player $i \in \text{Ag}$ we have a set of *actions* Ac_i and with every state s and player i we associate a subset $\text{Ac}_i(s) \subseteq \text{Ac}_i$ of actions that i can perform at s . We write Ac for $\bigcup_{i \in \text{Ag}} \text{Ac}_i$ and assume that the sets Ac_i form a partition of Ac . We call a profile of actions $(a_1, \dots, a_n) \in \text{Ac}_1 \times \dots \times \text{Ac}_n$ a *direction*, and denote it by d . We let D be the set of directions—also called *decisions*—with respect to Ac , and write d_i for the a_i of d that is in Ac_i . The dynamics of a game are modelled via a (deterministic) transition function $\delta: \text{St} \times \text{Ac}_1 \times \dots \times \text{Ac}_n \rightarrow \text{St}$, which indicates how the system behaves when $d = (a_1, \dots, a_n)$ is performed at a state s . A state s' is *accessible* from another state s whenever there is some $d = (a_1, \dots, a_n)$ such that $\delta(s, a_1, \dots, a_n) = s'$. A *run* is an infinite sequence $\rho = s_0 s_1 s_2 \dots$ such that for every $t \geq 0$ we have that s_{t+1} is accessible from s_t . The set of runs is denoted by R . By a (finite) *history* we mean a finite sequence $\pi = s_0 s_1 s_2 \dots s_k$ of accessible states. By $\text{prefix}(\rho)$ we denote the set of finite prefixes of ρ , i.e., $\text{prefix}(\rho) = \{\pi \in \text{St}^* : \rho = \pi \rho' \text{ for some } \rho' \in \text{St}^\omega\}$. Given $\pi \in \text{St}^+$, by $\text{last}(\pi)$ we denote the last state in π , i.e., if $\pi = \pi's$, then $\text{last}(\pi) = s$. By $s^0 \in \text{St}$ we denote the *initial state* associated with St .

A *strategy* for a player i is a function $f_i: \text{St}^+ \rightarrow \text{Ac}_i$ such that $f_i(\pi s) \in \text{Ac}_i(s)$ for every $\pi \in \text{St}^*$ and $s \in \text{St}$. That is, a strategy for a player i specifies for every finite history π an action available to i in $\text{last}(\pi)$. The set of strategies for player i is denoted by F_i . A *strategy profile* is a tuple $\vec{f} = (f_1, \dots, f_n)$ in $F_1 \times \dots \times F_n$. Observe that given a state s and a transition function $\delta: \text{St} \times \text{Ac}_1 \times \dots \times \text{Ac}_n \rightarrow \text{St}$, each strategy profile \vec{f} defines a unique run ρ where $\rho[0] = s$ and $\rho[t + 1] = \delta(\rho[t], f_1(\rho[0 \dots t]), \dots, f_n(\rho[0 \dots t]))$, for all $t \geq 0$. We write $\rho(\vec{f}, s)$ for such a run, and simply $\rho(\vec{f})$ if $s = s^0$. Furthermore, each player i has an associated dichotomous preference relation over runs, which is modelled as a subset Γ_i of the set of runs R . Intuitively, a player i strictly prefers all runs in Γ_i to those that are not in Γ_i and is indifferent otherwise. Thus, Γ_i represents the *goal* of player i . Below, we will use formulae of temporal logic to specify player's goals. We write $\rho \succsim_i \rho'$ to indicate that player i weakly prefers run ρ to ρ' and $\rho \succ_i \rho'$ for player i strictly preferring ρ to ρ' , i.e., if $\rho \succsim_i \rho'$ but not $\rho' \succsim_i \rho$. A game is played by each player i selecting a strategy f_i with the aim that the induced run $\rho(\vec{f})$ belongs to its goal set Γ_i . If $\rho(\vec{f}) \in \Gamma_i$ we say that i has its goal satisfied. Otherwise, we say that i does not have its goal satisfied.

Logics. Alternating-time Temporal Logic (ATL^* [1]) is an extension of CTL^* [9], a branching-time temporal logic, that allows for reasoning about games and strategies. More specifically, given a set of atomic propositions AP and a set of agents Ag , the language of ATL^* formulae is given by the following grammar:

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U} \phi \mid \langle\langle C \rangle\rangle \phi$$

such that $p \in \text{AP}$ and $C \subseteq \text{Ag}$, and the formulae \mathbf{X} and \mathbf{U} are in the scope of $\langle\langle C \rangle\rangle$, that is, are subformulae of a $\langle\langle C \rangle\rangle$ -formula. We use the following abbreviations: we write \top for $p \vee \neg p$, \perp for $\neg \top$, $\mathbf{F}\phi$ for $\top \mathbf{U} \phi$, $\mathbf{G}\phi$ for $\neg \mathbf{F} \neg \phi$, $\mathbf{E}\phi$ for $\langle\langle \text{Ag} \rangle\rangle \phi$, $\mathbf{A}\phi$ for $\langle\langle \emptyset \rangle\rangle \phi$, and $\llbracket C \rrbracket \phi$ for $\neg \langle\langle C \rangle\rangle \neg \phi$; we also use the conventional abbreviations for other classical propositional logic operators. We write $\phi \in \mathcal{L}(\text{AP}, \text{Ag})$ if ϕ is an ATL^* formula in this language. When either AP or Ag , or both, are known, we may omit them. With $\text{AP}' \subseteq \text{AP}$, we may write $\phi|_{\text{AP}'}$ if $\phi \in \mathcal{L}(\text{AP}', \text{Ag})$ for some set of agents Ag .

The semantics of ATL^* formulae are given by *concurrent game structures* [1]. A concurrent game structure, M , is given by a tuple $M = (\text{AP}, \text{Ag}, \text{Ac}, \text{St}, s^0, \lambda, \delta)$, where $\lambda: \text{St} \rightarrow 2^{\text{AP}}$ is a labelling function, and all other components of M are as defined before. The *size* of M is defined to be $|\text{St}| \times |\text{Ac}|^{|\text{Ag}|}$.

We write R_s^* and R_s^ω for, respectively, the finite and infinite runs of M that start at state s . We simply write R^* and R^ω if $s = s^0$. Moreover, we will write \vec{f}_C for (f_i, \dots, f_k) , with $C = \{i, \dots, k\} \subseteq \text{Ag}$, that is, a joint strategy for the players in C . Similarly, we will write \vec{g}_{-C} for a joint strategy for the players in Ag_{-C} . For simplicity, we will assume that all strategies are defined on all finite runs of M , and hence at all states. We define the set of \vec{f}_C -runs from state s to be $\{\rho' \in R_s^\omega : \rho' = \rho(\vec{f}_C(s), \vec{g}_{-C}(s)) \text{ for some } \vec{g}_{-C} \text{ for } \text{Ag}_{-C}\}$.

We can now define the semantics of ATL^* formulae based on the rules given below. Let M be a concurrent game structure, $\rho \in R^\omega$ be an infinite run, and $t \in \mathbb{N}$ be a temporal index. The semantics of ATL^* formulae is defined by the following rules:

$\rho, t \models_M p$	iff	$p \in \lambda(\rho[t])$
$\rho, t \models_M \neg\phi$	iff	$\rho, t \models_M \phi$ does not hold
$\rho, t \models_M \phi \vee \phi'$	iff	$\rho, t \models_M \phi$ or $\rho, t, \pi \models_M \phi'$
$\rho, t \models_M X\phi$	iff	$\rho, t+1 \models_M \phi$
$\rho, t \models_M \phi U \phi'$	iff	$\rho, t' \models_M \phi'$, for some $t' \geq t$, and $\rho, k \models_M \phi$, for all $t \leq k < t'$
$\rho, t \models_M \langle\langle C \rangle\rangle \phi$	iff	there is some \vec{f}_C such that for all \vec{f}_C -runs ρ' from state $\rho[t]$, it is the case that $\rho', 0 \models_M \phi$ holds.

We say that M is a *model* of ϕ (commonly written as $M \models \phi$) if $\rho, 0 \models_M \phi$ for all $\rho \in R^\omega$ such that $\rho[0] = s^0$. We also say that ϕ is *satisfiable* if $M \models \phi$ for some CGS M . Moreover, we say that ϕ is *equivalent* to ϕ' if $M \models \phi \iff M \models \phi'$ for all M , and define the *size* of ϕ as its number of subformulae. Finally, we define LTL as the sublogic of ATL* given by all formulae $A\phi$, where formula ϕ does not contain the “coalition” quantifiers $\langle\langle C \rangle\rangle$ or $\llbracket C \rrbracket$.

LTL Concurrent Games. An *LTL concurrent game* is given by a tuple $G = (M, \gamma_1, \dots, \gamma_n)$, such that M is a concurrent game structure where $\text{Ag} = \{1, \dots, n\}$ in which γ_i , with $i \in \text{Ag}$, is the LTL goal of player i that defines i 's preference relation over runs. Such a set of runs, for each player i , is defined as follows:

$$\Gamma_i = \{\rho \in R^\omega \mid \rho, 0 \models_M \gamma_i\}.$$

An *outcome* of G is a strategy profile $\vec{f} = (f_1, \dots, f_n)$ in $F_1 \times \dots \times F_n$ for the set of all agents (the *grand coalition*). In an LTL concurrent game (an *LTL game* hereafter) we can identify a set of “winners” and a set of “losers” for each outcome \vec{f} . Let $W(\vec{f})$ denote the set of players that would get their goal achieved if the outcome \vec{f} resulted (the “winners”) and let $L(\vec{f})$ denote the set of players that not:

$$\begin{aligned} W(\vec{f}) &= \{i \in \text{Ag} \mid \rho(\vec{f}) \models \gamma_i\} \\ L(\vec{f}) &= \text{Ag} \setminus W(\vec{f}). \end{aligned}$$

Before proceeding, it is useful to define the noncooperative solution concept of Nash equilibrium with respect to LTL concurrent games. An outcome \vec{f} is said to be a Nash equilibrium if there is no player $i \in \text{Ag}$ and strategy f'_i for i such that $i \in L(\vec{f})$ and $i \in W(\vec{f}_{-i}, f'_i)$. That is, \vec{f} is a Nash equilibrium if no player can benefit by unilaterally changing its strategy component of \vec{f} . Let $NE(G)$ denote the Nash equilibria of G . We emphasise that Nash equilibrium only considers *unilateral* deviations, *i.e.*, deviations by individual players. Our aim in what follows is to consider deviations that admit coalitions of players, in a cooperative setting.¹

3 COOPERATIVE RATIONAL VERIFICATION

Defining the Core. We want to define counterparts of the rational verification problems E-NASH and A-NASH, as studied in [13, 14], but for cooperative settings. For this, we need a version of the core for our cooperative game settings. The core is probably the best-known solution concept in cooperative game theory. Like Nash equilibrium in the noncooperative setting, the core defines a notion of stability for games, but whereas Nash equilibrium only requires that an outcome is stable in the sense that it admits no *individual* beneficial deviations, the core requires that an outcome admits no

beneficial deviations by *coalitions*. In the “standard” model of cooperative games, this intuition is easily formalised, but in concurrent games, there is an important difficulty. Suppose a coalition of players $C^* \subseteq \text{Ag}$ are contemplating an outcome \vec{f} , and in particular, are attempting to determine whether they have a cooperative beneficial deviation from \vec{f} . Now, as they consider possible beneficial deviations – collective strategies \vec{f}_{C^*} – *what assumptions should C^* make about the behaviour of the remaining players $\text{Ag} \setminus C^*$?* In particular, assuming that the remaining players will not alter their strategy is implausible in a cooperative setting:² rational players who can cooperate will respond to the deviation rationally and in a cooperative way against the players in C^* . And, crucially, whether or not C^* 's putative deviation is in fact beneficial may well depend upon the behaviour of the remaining players. In game theoretic terms, our concurrent game setting is subject to externalities: the performance of the coalition C^* depends not just on the coalition C^* , but on the behaviour of the remaining players.

It is well-known that cooperative solution concepts are difficult to define in the presence of externalities [4]. In particular, there is no universally accepted definition of the core for games with externalities. Our first definition of the core for concurrent games, therefore, captures *worst case* reasoning. Thus, when coalition C is contemplating a deviation, it requires that this deviation will be beneficial *no matter what the remaining players do*. This idea has been explored in the concept of the α -core in cooperative games [28]. To make this idea formal, we need to define the notion of a *deviation* and a *beneficial deviation*. A *deviation* is a joint strategy \vec{f}_{C^*} for the set of players $C^* \subseteq \text{Ag}$, with $C^* \neq \emptyset$. Where \vec{f} is an outcome, we say \vec{f}'_{C^*} is a *beneficial deviation* from \vec{f} if:

- (1) $C^* \subseteq L(\vec{f})$
- (2) For all \vec{f}'_{-C^*} , we have $C^* \subseteq W((\vec{f}'_{C^*}, \vec{f}'_{-C^*}))$.

In other words, \vec{f}'_{C^*} is said to be a beneficial deviation from \vec{f} if the players in C^* would be better off choosing strategies \vec{f}'_{C^*} , rather than their part of \vec{f} , no matter what strategies the players outside C^* chose. The core of a game G , denoted $\text{core}(G)$, is then defined to be the set of outcomes of G that admit no beneficial deviation.

Example 3.1. Consider the following game, which contains a poor quality Nash equilibrium that is not in the core: the ability to cooperate makes it possible for agents to avoid the undesirable equilibrium. The game contains two players, $\text{Ag} = \{1, 2\}$ and two variables $\text{AP} = \{p, q\}$, with player 1's action set being $\text{Ac}_1 = \{pt, pf\}$ and player 2's action set being $\text{Ac}_2 = \{qt, qf\}$, satisfying that, for every reachable state, if player 1/2 plays pt/qt then p/q will hold, and will not hold if pf/qf is played instead (*i.e.*, player 1 “controls” the value of p and player 2 the value of q). Their goals are identical (and so the game is a coordination game): $\gamma_1 = \gamma_2 = G(p \wedge q)$. Now, consider the strategy profile \vec{f} in which both players simply fix their respective variables to be false forever (*i.e.*, play pf and qf forever). Neither player will have their goal achieved by such a strategy profile. However, the strategy profile forms a Nash equilibrium, because unilateral deviation cannot improve the situation: neither player has an alternative strategy which would make

¹We recall that the concept strong Nash equilibrium also admits deviations by groups of players, but in a noncooperative setting.

²This is the kind of behaviour that one has to assume to define strong Nash equilibrium, a noncooperative solution concept.

them better off. In fact, there are infinitely many such poor quality Nash equilibria in this game, where neither player gets their goal achieved. However, this strategy profile is *not* in the core, because there is a cooperative beneficial deviation to the strategy profile in which both players fix their variables to be true forever (*i.e.*, play pt and qt forever). And, in fact, in every core-stable outcome, both players get their goal achieved. Thus, using the core instead of Nash equilibrium eliminates poor quality equilibria from the game, leading to socially more desirable outcomes.

Decision Problems. In Rational Verification [13, 14, 29] we are mainly interested in checking which temporal logic properties a game satisfies by its stable outcomes. Typically, in the noncooperative setting, such outcomes have been characterised by the set of Nash equilibria $NE(G)$ of the game G . In the cooperative setting, as introduced here, such outcomes are characterised, instead, by the set of outcomes in the core of the game, that is, by the strategy profiles in $core(G)$. The two main decision problems in rational verification are checking whether a temporal logic formula is satisfied by some/every stable outcome of the game. For the core, these problems are defined as follows—*cf.* [13, 14, 29].

E-CORE:

Given: Game G , LTL formula ϕ .

Question: Does $\exists \vec{f} \in core(G). \rho(\vec{f}) \models \phi$ hold?

A-CORE:

Given: Game G , LTL formula ϕ .

Question: Does $\forall \vec{f} \in core(G). \rho(\vec{f}) \models \phi$ hold?

In addition to the two above decision problems, the third main decision problem in rational verification is checking whether, given a game G , its set of stable outcomes—the core of G in this case—is non-empty. As will be shown in the next sections, in our setting, the core of every game G is never empty, a desirable game-theoretic property as it ensures the existence of stable outcomes for every game, making them rationally implementable in practice.

We will also be interested in two decision problems specifically related with the nature *core-stable outcomes* (that is, of outcomes in the core of a game), namely, checking whether a given outcome—a strategy profile for the grand coalition—is in the core (CORE MEMBERSHIP), and checking whether a given deviation is beneficial with respect to a given outcome of a game (BENEFICIAL DEVIATION). These two decision problems are formally defined as follows.

CORE MEMBERSHIP:

Given: Game G , outcome \vec{f} .

Question: Is it the case that $\vec{f} \in core(G)$?

BENEFICIAL DEVIATION:

Given: Game G , outcome \vec{f} , deviation \vec{f}'_{C^*} .

Question: Is \vec{f}'_{C^*} a beneficial deviation from \vec{f} ?

One might think that every coalition that has a beneficial deviation from some outcome of the game will get their goals achieved in a core-stable outcome, but that actually is not the case. To formalise this idea, let us introduce the concept of *fulfilled coalition*. We say that a coalition of players is fulfilled if they are able to achieve their goals irrespective of what other players do. Formally, we say that a coalition of players C is fulfilled if there is a joint strategy \vec{f}_C

for $C \subseteq Ag$ such that for all joint strategies \vec{f}_{-C} for $Ag \setminus C$ we have

$$\rho((\vec{f}_C, \vec{f}_{-C})) \models \bigwedge_{i \in C} \gamma_i.$$

In other words, a fulfilled coalition has a winning strategy to collectively achieve their goals. Since we are considering cooperative games, the issue/question is whether such a coalition will form. Using the above definition, we can make some useful observations about (fulfilled) coalitions and core-stable outcomes. These observations are formally presented in the following lemma, which relates winning strategies and the core in a critical way.

LEMMA 3.2 (COALITIONS).

- (1) *There are games G , with outcomes $\vec{f} \in core(G)$, containing fulfilled coalitions $C \subseteq Ag$ such that $C \not\subseteq W(\vec{f})$.*
- (2) *For every game G , outcome $\vec{f} \in core(G)$, and fulfilled coalition C , we have that $C \cap W(\vec{f}) \neq \emptyset$.*
- (3) *For every game G and fulfilled coalition C , if $core(G) \neq \emptyset$, then there is $\vec{f} \in core(G)$ such that $C \subseteq W(\vec{f})$.*

Informally, the first part of the lemma says that the fact that a coalition is fulfilled does not mean that every player in such a coalition is guaranteed to get its goal achieved in an arbitrary core-stable outcome. The second part of the lemma says that, however, in any core-stable outcome, some members of every fulfilled coalition must get their goals achieved. And, the third part of the lemma says that for every fulfilled coalition the core contains a core-stable outcome in which every member of this coalition gets its goal achieved. Because fulfilled coalitions can help us understand the coalition formation power in a game, we will also be interested in the following decision problem about coalitions.

FULFILLED COALITION:

Given: Game G , coalition $C \subseteq Ag$.

Question: Is C a fulfilled coalition of G ?

In the next section, we will investigate the decision problems defined here as well as some model-theoretic properties of the core.

4 REASONING ABOUT THE CORE

In this section we will study the computational complexity of the decision problems defined in the previous section, and will show some other properties of the core of an LTL game, in particular, that such a set is never empty and that the satisfaction of an LTL property on some/every outcome in the core is a bisimulation-invariant property [16]. These two results sharply contrast with the Rational Verification problem for noncooperative games in which, with respect to stable outcomes given by the set of Nash equilibria of a game (also given by strategy profiles for Ag), neither such a set of stable outcomes is guaranteed to always be non-empty [13] nor bisimulation-invariance holds in the general case [12].

The first decision problem we will consider in this section is FULFILLED COALITION, which we solve in the general case through a logical characterisation using ATL*.

THEOREM 4.1. *FULFILLED COALITION is PSPACE-complete for one-player games, and it is 2EXPTIME-complete for games with more than one player.*

PROOF. For membership we observe that given a game $G = (M, \gamma_1, \dots, \gamma_n)$ and a coalition $C \subseteq \text{Ag}$, it is the case that C is fulfilled if and only if $M \models \langle\langle C \rangle\rangle \bigwedge_{i \in C} \gamma_i$ holds [1]. Since such a formula can be model checked in PSPACE if Ag is a singleton set and in 2EXPTIME if $|C| > 1$, then the two upper bounds immediately follow. For the lower bounds, we can reduce the problem of checking for the existence of a winning strategy in a two-player game with LTL goals as defined in [2] for 2EXPTIME-hardness and existential LTL model checking for PSPACE-hardness [27]. \square

Fulfilled coalitions give an indication of which stable coalitions may form, but are insufficient to characterise the core, and therefore, to check E-CORE and A-CORE properties of a multi-agent system. To do this, we follow a different strategy and show that these two decision problems are, in general, also 2EXPTIME-complete.

THEOREM 4.2. *E-CORE and A-CORE are PSPACE-complete for one-player games and 2EXPTIME-complete for games with more than one player.*

PROOF. Let us consider E-CORE first. For membership we observe that given a game $G = (M, \gamma_1, \dots, \gamma_n)$ and an LTL formula ϕ , it is the case that $(G, \phi) \in \text{E-CORE}$ if and only if $M \models \phi_{\text{E-CORE}}(G, \phi)$ holds, such that $\phi_{\text{E-CORE}}(G, \phi)$ is the following ATL* formula:

$$\bigvee_{W \subseteq \text{Ag}} (\langle\langle \text{Ag} \rangle\rangle (\phi \wedge \bigwedge_{i \in W} \gamma_i \wedge \bigwedge_{j \in \text{Ag} \setminus W} \neg \gamma_j) \wedge \bigwedge_{L \subseteq \text{Ag} \setminus W} \llbracket L \rrbracket \bigvee_{j \in L} \neg \gamma_j)$$

which states that there is a path in M that satisfies ϕ and the goals of a set of players W (the “winners”) – subformula $\langle\langle \text{Ag} \rangle\rangle (\phi \wedge \bigwedge_{i \in W} \gamma_i \wedge \bigwedge_{j \in \text{Ag} \setminus W} \neg \gamma_j) \wedge \dots$ – and that for every subset of players L that do not get their goals achieved in such a path (the “losers”), it is not the case that those players have a beneficial deviation from the path – subformula $\langle\langle \text{Ag} \rangle\rangle (\dots) \wedge \bigwedge_{L \subseteq \text{Ag} \setminus W} \llbracket L \rrbracket \bigvee_{j \in L} \neg \gamma_j$. As in the case of fulfilled coalitions, when Ag is a singleton set, the question becomes an LTL model checking problem, which can be solved in PSPACE, and otherwise it is an ATL* model checking problem, which can be solved in 2EXPTIME, with the representation size of the game being exponential in the number of agents. For the lower bounds, it is sufficient to show that $(G, \phi) \in \text{E-CORE}$ if and only if $(G, \{1\}) \in \text{FULFILLED COALITION}$, whenever $\gamma_1 = \phi$ and $\gamma_j = \neg \phi$, for every $j \in \text{Ag} \setminus \{1\}$, which can be proved using Lemma 3.2. Since this is true even if Ag is a singleton set, both lower bounds follow from Theorem 4.1, that is, PSPACE-hardness in case of one-player games, and 2EXPTIME-hardness even for two-player games.

Finally, for A-CORE, we observe that $(G, \phi) \notin \text{A-CORE}$ if and only if $(G, \neg \phi) \in \text{E-CORE}$. Since both PSPACE and 2EXPTIME are deterministic complexity classes, we can conclude that A-CORE is PSPACE-complete if $|\text{Ag}| = 1$ and 2EXPTIME-complete if $|\text{Ag}| > 1$, as it is for E-CORE. \square

We now study CORE MEMBERSHIP and BENEFICIAL DEVIATION. For these two problems we first need to define how we will represent outcomes, as at present they are defined as infinite-state objects that map finite histories to players’ actions. Following standard practice in the concurrent games literature, we model strategies as finite state machines with output (transducers) [13, 14]. Note that, for players with LTL goals, such strategies are sufficient: no more powerful model of strategies is necessary [13, 14]. Formally, a strategy for player i is a structure $\sigma_i = (Q_i, q_i^0, \iota_i, \nu_i)$, where

- Q_i is a finite and non-empty set of *strategy states*,
- $q_i^0 \in Q_i$ is the *initial strategy state*,
- $\iota_i : Q_i \times \text{St} \rightarrow Q_i$ is a *transition function*, and
- $\nu_i : Q_i \rightarrow \text{Ac}_i$ is an *output function*, which satisfies that, for all $(q, s, q') \in \iota_i$, we have $\nu_i(q) \in \text{Ac}_i(s)$.

With this definition in place, we can now establish the complexity of CORE MEMBERSHIP and BENEFICIAL DEVIATION. Formally, we have the following result.

THEOREM 4.3. *CORE MEMBERSHIP is PSPACE-complete for one-player games and 2EXPTIME-complete for games with more than one player.*

PROOF. For membership we first compute the winners and losers with respect to $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$, the outcome of the game. This can be done in PSPACE (it is equivalent to LTL model checking over a “product automata” or “concurrent program” [18]). Once we have computed W , we can check, for every $L \subseteq \text{Ag} \setminus W$, whether L has a beneficial deviation. This is true if and only if L is a fulfilled coalition. Because this can be checked in PSPACE for one-player games and in 2EXPTIME for games with more than one player, the two upper bounds immediately follow. For the lower bounds, we use Lemma 3.2 and Theorem 4.1 again. Consider the following game. Let ϕ be a satisfiable LTL formula and $\vec{\sigma}$ an outcome that does not satisfy ϕ . Then, $(G, \vec{\sigma}) \in \text{CORE MEMBERSHIP}$ if and only if $(G, \{1\}) \notin \text{FULFILLED COALITION}$, whenever $\gamma_1 = \phi$ and $\gamma_j = \neg \phi$, for every player $j \in \text{Ag} \setminus \{1\}$. \square

Let us now consider BENEFICIAL DEVIATION. This is the only “easy” problem for multi-player games, as it can be solved in PSPACE. To show this, we again need to find a different proof strategy. Consider any input instance $(G, \vec{\sigma}, \vec{\sigma}'_{C^*})$ of the problem. We observe that, because $\vec{\sigma}'_{C^*}$ is fixed, we can make it part of the arena where the game is played, and then check if players not in C^* have a joint strategy for $\bigvee_{j \in C^*} \neg \gamma_j$. Due to the definition of beneficial deviation, we also need to check if $\rho(\vec{\sigma}) \models \bigwedge_{j \in C^*} \neg \gamma_j$ holds or not.

In other words, the reason why this problem can be solved in PSPACE for multi-player games, unlike all other decision problems we have studied so far (which, in general, can be solved in doubly exponential time), is that this decision problem can be reduced to a one-player game (given by coalition $\text{Ag} \setminus C^*$) with an LTL goal (given by $\gamma_{\text{Ag} \setminus C^*} = \bigvee_{j \in C^*} \neg \gamma_j$) over a “product arena” (denoted by M_{C^*}) built from a concurrent game structure M and the joint strategy \vec{f}'_{C^*} that we want to check.

THEOREM 4.4. *BENEFICIAL DEVIATION is PSPACE-complete, even for one-player games.*

PROOF. Checking that $\rho(\vec{\sigma}) \models \bigwedge_{j \in C^*} \neg \gamma_j$ holds can be done in PSPACE. Again, this is equivalent to model checking LTL formulae over a “product automata” or “concurrent program” [18]). If the statement does not hold, then, by definition, $\vec{\sigma}'_{C^*}$ is not a beneficial deviation, as at least one player in C^* already has its goal satisfied by $\vec{\sigma}$. If the statement holds, then we check that $\rho(\vec{f}'_{C^*}, \vec{\sigma}'_{C^*}) \models \bigwedge_{j \in C^*} \neg \gamma_j$ holds, for all joint strategies \vec{f}'_{C^*} for players not in C^* . We do this in PSPACE by checking whether it is not the case that $M_{C^*} \models \bigvee_{j \in C^*} \neg \gamma_j$ holds, where $M_{C^*} = (\text{AP}', \text{Ag}', \text{Ac}', \text{St}', s^0, \lambda', \delta')$ is the “concurrent program” or “product automata” defined as follows:

- $AP' = AP$, $Ag' = \{0\}$, $Ac' = \Pi_{i \in Ag \setminus C^*} Ac_i$;
- $St' = St \times \Pi_{j \in C^*} Q_j$;
- $s^{0'} = (s^0, q_x^0, \dots, q_y^0)$, such that $\sigma_z = (Q_z, q_z^0, \iota_z, v_z)$, $\vec{\sigma}_{C^*} = (\sigma_x, \dots, \sigma_y)$, and $z \in \{x, \dots, y\}$;
- $\lambda'(s, q_x, \dots, q_y) = \lambda(s)$; and
- $\delta'((s, q_x, \dots, q_y), (a, \dots, b)) = (s', q'_x, \dots, q'_y)$ such that
 - $s' = \delta(s, v(q_x), \dots, v(q_y), a, \dots, b)$, and
 - $q'_z = \iota(q_z, s)$, with $z \in \{x, \dots, y\}$.

In other words, M_{C^*} transitions just like M save that it is restricted to the behaviour already defined by $\vec{\sigma}_{C^*}$.

For the lower bound we use LTL model checking. \square

In addition to the above complexity results, we also have two model-theoretic results, one ensuring that the core is never empty and another one stating that checking whether an LTL formula is satisfied by some outcome in the core is a bisimulation-invariant property.³ The latter result is easy, and follows directly from the membership proof of E-CORE.

COROLLARY 4.5. *Let $G = (M, \gamma_1, \dots, \gamma_n)$ be a game, ϕ be an LTL formula, and M' be a concurrent game structure that is bisimilar to M . Then, $(G, \phi) \in E\text{-CORE}$ if and only if $(G', \phi) \in E\text{-CORE}$, where $G' = (M', \gamma_1, \dots, \gamma_n)$.*

PROOF. Because ATL^* is a bisimulation-invariant temporal logic, and the core can be characterised in ATL^* using $\phi_{E\text{-CORE}}$, as defined in the membership proof of E-CORE. More specifically, it follows from the fact that $M \models \phi_{E\text{-CORE}}(G, \phi)$ if and only if $M' \models \phi_{E\text{-CORE}}(G', \phi)$. \square

To finish this section, we show an important property of the core, namely, that it is never empty.

THEOREM 4.6. *$core(G) \neq \emptyset$, for every game G .*

PROOF. Take any run ρ in the game $G = (M, \gamma_1, \dots, \gamma_n)$. Either $\rho, 0 \models_M \bigwedge_{i \in Ag} \gamma_i$ holds or not. If the former, then the core is not empty: every outcome \vec{f} such that $\rho = \rho(\vec{f})$ is in the core. If the latter, then there is a set of players L_1 that do not get their goals achieved in ρ . If no subset of L_1 is fulfilled, then, again, every outcome \vec{f} such that $\rho = \rho(\vec{f})$ is in the core, since no set of losers would be able to beneficially deviate. Otherwise, there is a set of players $C_1 \subseteq L_1$ that have a joint strategy \vec{f}_{C_1} such that $\rho(\vec{f}_{C_1}, \vec{f}'_{-C_1}) \models \bigwedge_{i \in C_1} \gamma_i$, for all joint strategies \vec{f}'_{-C_1} for $Ag \setminus C_1$.

Now, take any outcome $(\vec{f}_{C_1}, \vec{f}'_{-C_1})$, that is, any outcome such that $C_1 \subseteq W(\vec{f}_{C_1}, \vec{f}'_{-C_1})$. Let $L_2 = L(\vec{f}_{C_1}, \vec{f}'_{-C_1})$. Again, if no subset of L_2 is fulfilled, then $(\vec{f}_{C_1}, \vec{f}'_{-C_1}) \in core(G)$. Otherwise, there are players $C_2 \subseteq L_2$ that have a joint strategy \vec{f}_{C_2} such that

$$\rho(\vec{f}_{C_1}, \vec{f}_{C_2}, \vec{f}'_{-(C_1 \cup C_2)}) \models \bigwedge_{i \in C_1 \cup C_2} \gamma_i$$

for all joint strategies $\vec{f}'_{-(C_1 \cup C_2)}$ for $Ag \setminus (C_1 \cup C_2)$.

We can now reason recursively and take this time any outcome $(\vec{f}_{C_1}, \vec{f}_{C_2}, \vec{f}'_{-(C_1 \cup C_2)})$, that is, any outcome such that $(C_1 \cup C_2) \subseteq W(\vec{f}_{C_1}, \vec{f}_{C_2}, \vec{f}'_{-(C_1 \cup C_2)})$, and let $L_3 = L(\vec{f}_{C_1}, \vec{f}_{C_2}, \vec{f}'_{-(C_1 \cup C_2)})$. Same

reasoning above applies, and since $L_1 \supset L_2 \supset L_3 \supset \dots$ is at most of length $|Ag|$, we know that either L_k , with $1 \leq k \leq |Ag|$, is not empty and an element in the core was found, essentially, any outcome $(\vec{f}_{C_1}, \vec{f}_{C_2}, \dots, \vec{f}_{C_{k-1}}, \vec{f}'_{-\bigcup_{1 \leq i < k} C_i})$ or L_k is empty, in whose case the strategy profile $(\vec{f}_{C_1}, \vec{f}_{C_2}, \dots, \vec{f}_{C_k})$ is necessarily in the core, which concludes the proof. Note that while any outcome in the core of the form $(\vec{f}_{C_1}, \vec{f}_{C_2}, \dots, \vec{f}_{C_{k-1}}, \vec{f}'_{-\bigcup_{1 \leq i < k} C_i})$ satisfies the goals of all players in $\bigcup_{1 \leq i < k} C_i$, any outcome $(\vec{f}_{C_1}, \vec{f}_{C_2}, \dots, \vec{f}_{C_k})$ in the core satisfies all of players' goals. \square

Theorem 4.6, ensuring that the core is never empty, can be used to strengthen numeral 3 of Lemma 3.2.

COROLLARY 4.7. *For every game G and fulfilled coalition C , there is $\vec{f} \in core(G)$ such that $C \subseteq W(\vec{f})$.*

5 ON CREDIBLE COALITION FORMATION

As we noted above, our definition of the core assumes worst-case reasoning: a deviation must be beneficial against *all* counterpart behaviours. This definition is robust in the sense that any core-stable outcome is stable in a very strong sense, but one could argue that in some cases it is *too* strong. In particular, when a coalition C^* is contemplating a deviation \vec{f}_{C^*} , it can surely assume that the remaining players will not act against their own interests. Thus, one could argue that a deviation need not be beneficial for *all* behaviours of the remaining players, but only those behaviours that are *credible*, in the sense that the remaining players might rationally choose them according to their own preferences.

To make this discussion concrete, consider a two-player game G containing only three infinite runs (see Figure 1), $\rho_\emptyset, \rho_{\{1\}}, \rho_{\{2\}}$, and four outcomes $\vec{f}_{a_1 a_2}, \vec{f}_{a_1 b_2}, \vec{f}_{b_1 a_2}, \vec{f}_{b_1 b_2}$, where $\rho(\vec{f}_{a_1 a_2}) = \rho_{\{1\}}$, $\rho(\vec{f}_{a_1 b_2}) = \rho_{\{2\}}$, $\rho(\vec{f}_{b_1 a_2}) = \rho_\emptyset = \rho(\vec{f}_{b_1 b_2})$, such that only $\rho_{\{1\}}$ satisfies γ_1 and only $\rho_{\{2\}}$ satisfies γ_2 . Now, in this game, $\vec{f}_{a_1 a_2} \in core(G)$ holds, but with the use of an incredible (punishing) strategy by player 1. Notice that the only possible deviation from $\vec{f}_{a_1 a_2}$ for player 2 is to $\vec{f}_{a_1 b_2}$ and hence the only possible response for player 1 is to $\vec{f}_{b_1 b_2}$. Although this behaviour would prevent player 2 from achieving its goal, such a way of playing can be regarded as not rational for player 1 given his preference relation: player 1 certainly prefers $\rho_{\{1\}}$ over the other two runs, but he is indifferent otherwise.




Figure 1: A game with an incredible strategy.

Motivated by this phenomenon, we propose a stronger definition for the core in which the way that deviating players are punished

³The reader is referred to [12] for the definition of bisimulation-invariance over the model of concurrent game structures.

is more credible. More specifically, with this new definition we require that if a coalition C^* wants to deviate from a given outcome \vec{f} using a joint strategy \vec{f}'_{C^*} , the coalition of players outside C^* can credibly threaten C^* only if players outside C^* have a joint strategy \vec{f}'_{-C^*} with which both at least one player in C^* does not get its goal achieved and every winner in \vec{f} remains a winner in $(\vec{f}'_{-C^*}, \vec{f}'_{C^*})$, that is, they act in accordance with their preference relations. We then reformulate the definition of a beneficial deviation and say that a deviation \vec{f}'_{C^*} is a *beneficial deviation* from \vec{f} if:

- (1) $C^* \subseteq L(\vec{f})$, and
- (2) $C^* \subseteq W(\vec{f}_{-C^*}, \vec{f}'_{C^*})$, and
- (3) for every joint strategy \vec{f}'_{-C^*} for $\text{Ag} \setminus C^*$ we have $W(\vec{f}) \subseteq W(\vec{f}'_{-C^*}, \vec{f}'_{C^*}) \Rightarrow C^* \subseteq W(\vec{f}'_{-C^*}, \vec{f}'_{C^*})$.

With this definition in place we can say that the *strong core* of a game (CORE^+), denoted $\text{core}^+(G)$, is the set of outcomes of G that admit no beneficial deviation as above. Then, we see that while $\vec{f}_{a_1 a_1}$ is in $\text{core}(G)$, it is not the case that $\vec{f}_{a_1 a_1}$ is in $\text{core}^+(G)$, since player 2 can now beneficially deviate from $\vec{f}_{a_1 a_1}$ to $\vec{f}_{a_1 b_1}$.

Note on credible threats in games with externalities: The game theory literature on this topic is vast. The reason is that the existence of externalities leads to many different definitions of stable behaviour (see, e.g., [10, 23, 28, 30] for many variants of the core). Here, we propose one definition but by no means we claim it is the strongest anyone may wish to consider. Essentially, with our definition, we require that for a punishing joint strategy to be credible, winners must remain winners after the presenting the threat.

We will now study the complexity of the decision problems defined in previous sections, but with respect to CORE^+ . There are four decision problems whose definition depends on the nature of the core: E-CORE, A-CORE, CORE MEMBERSHIP, and BENEFICIAL DEVIATION. To simplify notations, we will call them here in the same way but with the understanding that results in this section are with respect to CORE^+ . As we will show next, these four problems have the same complexities as with core, but require a more complex logical characterisation, which we provide here using the alternation fragment of Strategy Logic (SL [21, 22]).⁴

SL extends LTL with two *strategy quantifiers*, $\langle\langle x \rangle\rangle$ and $[[x]]$, and an *agent binding* operator (i, x) , where i is an agent and x is a variable. These operators can be read as “there exists a strategy x ”, “for every strategy x ”, and “bind agent i to the strategy associated with variable x ”, respectively. Formally, SL formulae are inductively built from a set of propositions AP, variables Var, and agents Ag, using the following grammar, where $p \in \text{AP}$, $x \in \text{Var}$, and $i \in \text{Ag}$:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U}\phi \mid \langle\langle x \rangle\rangle\phi \mid [[x]]\phi \mid (i, x)\phi.$$

We can now present the semantics of SL, where Str denotes the set of all strategies. Given a concurrent game structure M , for all SL formulae ϕ , states $s \in \text{St}$ in M , and assignments $\chi \in \text{Asg} = (\text{Var} \cup \text{Ag}) \rightarrow \text{Str}$, mapping variables and agents to strategies, the relation $M, \chi, s \models \phi$ is defined as follows:

- (1) For the Boolean and temporal cases, the semantics is standard
- (2) For all formulae ϕ and variables $x \in \text{Var}$ we have:
 - (a) $M, \chi, s \models \langle\langle x \rangle\rangle\phi$ if $\exists f \in \text{Str}. M, \chi[x \mapsto f], s \models \phi$;

- (b) $M, \chi, s \models [[x]]\phi$ if $\forall f \in \text{Str}. M, \chi[x \mapsto f], s \models \phi$;
- (3) For every agent $i \in \text{Ag}$ and variable $x \in \text{Var}$,
 - if $M, \chi[i \mapsto \chi(x)], s \models \phi$ then $M, \chi, s \models (i, x)\phi$.

For a sentence ϕ , that is, a formula with no free variables and agents [21, 22], we say that M satisfies ϕ , and write $M \models \phi$ in that case, if $M, \emptyset, s^0 \models \phi$, where \emptyset is the empty assignment. We use the following abbreviations: $\langle i \rangle\phi$ for $\langle\langle x \rangle\rangle(i, x)\phi$ and $[i]\phi$ for $[[x]](i, x)\phi$, which can be intuitively understood as “there is a strategy for agent i such that ϕ holds” and “ ϕ holds, for all strategies of agent i ”, respectively. We extend this notation to sets of players and write, for instance, $\langle C \rangle\phi$ instead of $\langle i \rangle \dots \langle j \rangle\phi$, where $C = \{i, \dots, j\}$, and similarly for the universal quantifier operator. Then, with $\langle C \rangle\phi$ we mean that “coalition C has a joint strategy such that ϕ holds.”

We then find that for a game $G = (M, \gamma_1, \dots, \gamma_n)$ and LTL formula ϕ , we have $(G, \phi) \in \text{E-CORE}$ if and only if $M \models \phi_{\text{E-CORE}}^+(G, \phi)$, where $\phi_{\text{E-CORE}}^+(G, \phi)$ is the SL formula:

$$\begin{aligned} \phi_{\text{E-CORE}}^+(G, \phi) &= \bigvee_{W \subseteq \text{Ag}} \langle \text{Ag} \rangle (\\ &\quad \phi \wedge \bigwedge_{i \in W} \gamma_i \wedge \bigwedge_{j \in \text{Ag} \setminus W} \neg \gamma_j \wedge \\ &\quad \bigwedge_{C^* \subseteq \text{Ag} \setminus W} \phi_{\text{NoBD}}(G, W, C^*)) \\ \phi_{\text{NoBD}}(G, W, C^*) &= [C^*] (\bigwedge_{j \in C^*} \gamma_j \rightarrow \\ &\quad \langle \text{Ag} \setminus C^* \rangle (\bigwedge_{i \in W} \gamma_i \wedge \bigvee_{j \in C^*} \neg \gamma_j)) \end{aligned}$$

This SL formula expresses that in the concurrent game structure there exists a path $\langle\langle \text{Ag} \rangle\rangle \dots$ satisfying that formula ϕ holds, some players get their goals achieved ($\bigwedge_{i \in W} \gamma_i$) – the winners, the remaining players do not ($\bigwedge_{j \in \text{Ag} \setminus W} \neg \gamma_j$) – the losers, and no coalition of losers have a beneficial deviation ($\bigwedge_{C^* \subseteq \text{Ag} \setminus W} \phi_{\text{NoBD}}(G, W, C^*)$). In addition, a coalition of losers C^* not having a beneficial deviation is expressed with the SL formula $\phi_{\text{NoBD}}(G, W, C^*)$ as follows. For every joint strategy of C^* (formula $[C^*] \dots$), if with such a joint strategy every player in C^* is better off ($\bigwedge_{j \in C^*} \gamma_j$ – for condition (2) of beneficial deviation) then the coalition of players outside C^* have a joint strategy ($\langle \text{Ag} \setminus C^* \rangle \dots$) such that both the winners in the original outcome remain winners after the threat is presented ($\bigwedge_{i \in W} \gamma_i$), and at least one player in the deviating coalition, C^* , does not get its goal achieved ($\bigvee_{j \in C^*} \neg \gamma_j$) – for condition (3) of the definition of beneficial deviation with respect to CORE^+ .

At this point, we would like to make a couple of observations. First, that the complexity of checking SL formulae is non-elementary and depends on the alternation-depth of the formula ([21]): SL formulae of alternation-depth n can be checked in $(n + 1)$ -EXPTIME, and in PSPACE for formulae that are semantically equivalent to CTL* formulae. Since $\phi_{\text{E-CORE}}^+(G, \phi)$ is an SL formula with two alternations, it can be checked in 3EXPTIME (and in PSPACE if $|\text{Ag}| = 1$). Second, we also would like to recall that finite-state machine strategies, as those we use here, can be characterised in LTL using the technique presented in [13, 14]. Using these logical characterisations, we can show the following complexity results.

THEOREM 5.1. *For multi-player games, while E-CORE and A-CORE are in 3EXPTIME, CORE MEMBERSHIP is 2EXPTIME-complete and BENEFICIAL DEVIATION is PSPACE-complete. For one-player games, all problems are PSPACE-complete.*

Because CORE^+ was characterised using SL (which is not a bisimulation-invariant logic), we cannot conclude that the satisfaction of LTL properties by outcomes in core^+ is a bisimulation-invariant property. We believe that this is not the case. Furthermore,

⁴A logical characterisation of CORE^+ using ATL* was not found. In fact, we believe that such a logical characterisation in ATL* is not possible for multi-player games.

we do not know whether $core^+$ is never empty or not. As a partial result, we show that such a set is not empty in two-player games.

PROPOSITION 5.2. $core^+ \neq \emptyset$, for two-player games.

PROOF. For a contradiction, let us suppose that for some game G , the set of outcomes $core^+(G)$ is empty. This means that for every outcome either player 1 or player 2 or both have a beneficial deviation. Then, we know that no outcome can satisfy both goals, γ_1 and γ_2 . Let us then consider the three remaining possible cases: outcomes that only satisfy γ_1 (case 1), outcomes that only satisfy γ_2 (case 2), and outcomes that satisfy neither γ_1 nor γ_2 (case 3). Let $\vec{f} = (f_1, f_2)$ be an outcome, f'_1 be a deviation by player 1, and f'_2 be a deviation by player 2, and consider the three cases above. In case 1, only player 2 would deviate. Then, outcome (f_1, f'_2) only satisfies γ_2 . Because (f_1, f'_2) is not in the core either, from this outcome only player 1 would deviate, to another outcome (f'_1, f'_2) . Then, outcome (f'_1, f'_2) only satisfies γ_1 . But, then, we have a contradiction, since this means that (f_1, f_2) would be in $core^+(G)$. We can reason symmetrically to show that case 2 is not possible either. For case 3 we note that only single deviations would be possible. But any such deviations would be to an outcome that either only satisfies γ_1 or only satisfies γ_2 , which are no longer possible. Since no other cases are possible, we have to reject our assumption and conclude that, for two-player games, $core^+(G)$ is not empty. \square

6 DISCUSSION AND RELATED WORK

Coalition formation in cooperative games. Coalition formation with externalities has been studied in the cooperative game-theory literature [10, 28, 30]. They considered several concepts of the core. For instance, α -core takes the pessimistic approach that requires that all members of a deviating coalition, S , will benefit from the deviation regardless of the behaviour of the other coalitions that may be formed. Our first definition of the core follows this approach. Contrarily, β -core takes an optimistic approach and requires that the members of a deviating coalition S will benefit from at least one possibility of coalition formation of the rest of the players. In addition, γ -core [5, 6] assumes that the coalition structure that will be created after a deviation will include the deviating coalition S and the rest of the coalition structure will consist of all singletons. The “worth” of S is now defined as equal to its payoff in the Nash equilibrium between S and the other players acting individually, in which the members of S play their joint best response strategy against the individually best response strategies of the remaining players. It is well-known that α - and β -characteristic functions lead to large cores [25], which is consistent with our observation that, with respect to our first definition, the core is never empty. Coalition formation is important in multi-agent system [26]. However, even though Coalition formation with externalities is very common in multi-agent systems, not much work has studied the concept of stability in multi-agent coalition formation with externalities. Instead, in artificial intelligence and multi-agent systems, most research has focused on the structure formation itself [24].

Quantitative reasoning in concurrent games. In this paper, we studied games where players goals and preferences are defined using LTL formulae. This led to a very general qualitative analysis of their rational behaviour, in particular, regarding the coalitions

they may form. Another criteria to decide whether a coalition of players might form is with the use of quantitative information, for instance, when considering mean-payoff functions, or some other kind of utility functions instead. More specifically, to investigate this quantitative setting, one could associate states in a concurrent game structure with an n -tuple (v_1, \dots, v_n) of positive real values, one for each agent i in the game, through a *value* function $val(i) : St \rightarrow \mathbb{R}$. Games will be played in the same way, and hence produce infinite runs $\rho = s^0 s_1 s_2 \dots$ which each will induce an infinite sequence of values for each player: $val_i(\rho) = v_0 v_1 \dots$, where for each $i \in Ag$, we have a map $val_i : R^\omega \rightarrow \mathbb{R}^\omega$. The payoff of a player i , denoted by mp_i , on a given outcome \vec{f} will be the *mean-payoff* of the infinite sequence of values associated with the run induced by such an outcome, which is formally defined as follows:

$$mp_i(val_i(\rho(\vec{f}))) = \liminf_{k \rightarrow \infty} \frac{1}{k} \sum_{t=0}^k val_i[t].$$

Preference relations are defined in the obvious way. Player i strictly prefers outcome \vec{f} over outcome \vec{f}' if and only if $mp_i(\vec{f}) > mp_i(\vec{f}')$, and is indifferent otherwise. On this basis, for a set of players $C^* \subseteq Ag$, a deviation \vec{f}'_{C^*} is a *beneficial deviation from \vec{f}* if and only if, for each $i \in C^*$, we have $mp_i(\vec{f}'_{C^*}, \vec{f}'_{-C^*}) > mp_i(\vec{f})$, for every \vec{f}'_{-C^*} of $Ag \setminus C^*$. Then, the core of a game G with quantitative payoffs would be the set of outcomes of G that admit no beneficial deviation. With these definitions in place one can now ask the main questions studied before, namely, E-CORE, A-CORE, CORE MEMBERSHIP, and BENEFICIAL DEVIATION, but in an unexplored quantitative setting. We believe that without LTL goals associated with the players in the game, this setting may lead to better complexity results, but this is something that still has to be fully investigated.

Rational verification of concurrent games. The formal verification of temporal logic properties of multi-agent systems, while assuming rational behaviour of the agents in such a system, has been studied for almost a decade now; see, for instance, [11, 13, 14, 17, 29]. However, to the best of our knowledge, all these studies have considered a non-cooperative setting, even if coalitional power is allowed, for instance, as in a strong Nash equilibrium. Nonetheless, also in such non-cooperative settings, the complexity of checking whether a temporal logic property is satisfied in a stable outcome of the game is a 2EXPTIME-complete problem, even for two-player zero-sum games where only trivial coalitions can be formed. On the positive side, cooperative games seem to have better model theoretic properties in the rational verification framework: with respect to our first definition of the core (which corresponds to the concept of α -core in the literature of cooperative games), a witness in the core is always guaranteed (since the core is never empty), preserved across bisimilar systems, and easily checked in practice in an efficient way using ATL* model checking techniques, which are supported by a number of verification tools, e.g., MCMAS [19], an automated formal verification tool that supports ATL* model checking and specifications of concurrent game structures in ISPL.

Acknowledgment. This research has been partly supported by the Ministry of Science and Technology, Israel, and the Japan Science and Technology Agency (JST).

REFERENCES

- [1] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5 (2002), 672–713.
- [2] Rajeev Alur and Salvatore La Torre. 2004. Deterministic generators and games for LTL fragments. *ACM Transactions on Computational Logic* 5, 1 (2004), 1–25.
- [3] Romain Brenguier. 2013. PRALINE: A Tool for Computing Nash Equilibria in Concurrent Games. In *CAV (Lecture Notes in Computer Science)*, Vol. 8044. Springer, 890–895.
- [4] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan-Claypool.
- [5] Parkash Chander. 2007. The gamma-core and coalition formation. *International Journal of Game Theory* 35, 4 (2007), 539–556.
- [6] Parkash Chander et al. 2010. Cores of games with positive externalities. *CORE DP* 2010/4 (2010).
- [7] Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. 2002. *Model Checking*. MIT Press.
- [8] E. Allen Emerson. 1990. Temporal and Modal Logic. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*. MIT Press, 995–1072.
- [9] E. Allen Emerson and Joseph Y. Halpern. 1986. "Sometimes" and "Not Never" revisited: on branching versus linear time temporal logic. *Journal of the ACM* 33, 1 (1986), 151–178.
- [10] Michael Finus and Bianca Rundshagen. 2003. A Non-Cooperative Foundation of Core-Stability in Positive Externality NTU-Coalition Games. *Nota Di Lavoro* 31.2003. *Economics Energy Environment* (2003).
- [11] Dana Fisman, Orna Kupferman, and Yoav Lustig. 2010. Rational Synthesis. In *TACAS (Lecture Notes in Computer Science)*, Vol. 6015. Springer, 190–204.
- [12] Julian Gutierrez, Paul Harrenstein, Giuseppe Perelli, and Michael Wooldridge. 2017. Nash Equilibrium and Bisimulation Invariance. In *CONCUR (LIPIcs)*, Vol. 85. Schloss Dagstuhl, 17:1–17:16.
- [13] Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. 2015. Iterated Boolean games. *Information and Computation* 242 (2015), 53–79.
- [14] Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. 2017. From model checking to equilibrium checking: Reactive modules for rational verification. *Artificial Intelligence* 248 (2017), 123–157.
- [15] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael Wooldridge. 2018. EVE: A Tool for Temporal Equilibrium Analysis. In *ATVA (Lecture Notes in Computer Science)*, Vol. 11138. Springer, 551–557.
- [16] Matthew Hennessy and Robin Milner. 1985. Algebraic Laws for Nondeterminism and Concurrency. *Journal of the ACM* 32, 1 (1985), 137–161.
- [17] Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. 2016. Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence* 78, 1 (2016), 3–20.
- [18] Orna Kupferman, Moshe Y. Vardi, and Pierre Wolper. 2000. An automata-theoretic approach to branching-time model checking. *Journal of the ACM* 47, 2 (2000), 312–360.
- [19] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. 2017. MCMAS: an open-source model checker for the verification of multi-agent systems. *STTT* 19, 1 (2017), 9–30.
- [20] Michael Maschler, Eilon Solan, and Shmuel Zamir. 2013. *Game Theory*. Cambridge University Press.
- [21] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Transactions on Computational Logic* 15, 4 (2014), 34:1–34:47.
- [22] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2017. Reasoning about Strategies: on the Satisfiability Problem. *Logical Methods in Computer Science* 13, 1 (2017).
- [23] Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT Press.
- [24] Talal Rahwan, Tomasz Michalak, Michael Wooldridge, and Nicholas R Jennings. 2012. Anytime coalition structure generation in multi-agent systems with positive or negative externalities. *Artificial Intelligence* 186 (2012), 95–122.
- [25] Debraj Ray and Rajiv Vohra. 1997. Equilibrium binding agreements. *Journal of Economic theory* 73, 1 (1997), 30–78.
- [26] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial intelligence* 101, 1 (1998), 165–200.
- [27] A. Prasad Sistla and Edmund M. Clarke. 1985. The Complexity of Propositional Linear Temporal Logics. *Journal of the ACM* 32, 3 (1985), 733–749.
- [28] Metin Uyanik. 2015. On the nonemptiness of the α -core of discontinuous games: Transferable and nontransferable utilities. *Journal of Economic Theory* 158 (2015), 213–231.
- [29] Michael Wooldridge, Julian Gutierrez, Paul Harrenstein, Enrico Marchioni, Giuseppe Perelli, and Alexis Toumi. 2016. Rational Verification: From Model Checking to Equilibrium Checking. In *Proc. of AAAI*. AAAI Press, 4184–4191.
- [30] Sang-Seung Yi. 1997. Stable coalition structures with externalities. *Games and economic behavior* 20, 2 (1997), 201–237.