

Multi-Player Games with LDL Goals over Finite Traces

Julian Gutierrez¹, Giuseppe Perelli^{2†}, Michael Wooldridge¹

¹*Department of Computer Science, University of Oxford*

²*Department of Informatics, University of Leicester*

Abstract

Linear Dynamic Logic on finite traces (LDL_F) is a powerful logic for reasoning about the behaviour of concurrent and multi-agent systems. In this paper, we investigate techniques for both the *characterisation* and *verification* of equilibria in multi-player games with goals/objectives expressed using logics based on LDL_F . This study builds upon a generalisation of *Boolean games*, a logic-based game model of multi-agent systems where players have goals succinctly represented in a logical way. Because LDL_F goals are considered, in the settings we study—Reactive Modules games and iterated Boolean games with goals over finite traces—players’ goals can be defined to be regular properties while achieved in a finite, but arbitrarily large, trace. In particular, using *alternating automata*, the paper investigates automata-theoretic approaches to the characterisation and verification of (pure strategy Nash) equilibria, shows that the set of Nash equilibria in multi-player games with LDL_F objectives is regular, and provides complexity results for the associated automata constructions.

Keywords: Games, Temporal Logic, Multi-Agent Systems, Formal Verification.

1. Introduction

Boolean games (BG [?]) are a logic-based model of multi-agent systems where each agent/player i is associated with a goal, represented as a *propositional logic* (PL) formula γ_i , and player i ’s main purpose is to ensure that γ_i is satisfied. The strategies and choices for each player i are defined with respect to a set of Boolean variables Φ_i , drawn from an overall set of variables Φ . Player i is assumed to have unique control over the variables in Φ_i , in that it can assign truth values to these variables in any way it chooses. Strategic concerns arise in Boolean games as the satisfaction of player i ’s goal γ_i can depend on the variables controlled by other players.

Reactive Modules games (RMG [?]) and *iterated Boolean games* (iBG [?]) generalise Boolean games by making players interact with each other for infinitely many rounds. As in the standard (one-shot or one-round) setting described above, there are n players each of whom uniquely controls a subset of Boolean variables and defines the achievement of a particular goal formula γ_i satisfied. However, in RMGs

[†]This work has been done while being affiliated to the University of Oxford.

and iBGs, players' goals γ_i are *Linear Temporal Logic* formulae (LTL [?]), rather than PL formulae, which are naturally interpreted over infinite sequences of valuations of the variables in Φ ; thus, in both RMGs and iBGs, such infinite sequences of valuations represent the plays of these games.¹

Even though RMGs, iBGs, and conventional Boolean games are logic-based models of multi-agent systems, they capture players' goals—and therefore the desired behaviour of the underlying multi-agent systems—in radically different ways: whereas Boolean games have PL goals (which are naturally evaluated over one-round games), RMGs and iBGs have LTL goals (which are naturally evaluated over games with infinitely many rounds), encompassing two extremes of the landscape when considering repeated games. However, there are games, systems, or situations where goals evaluated after an unbounded, but certainly finite, number of rounds should, or must, be considered [? ?].

In this paper we fill this gap and define and investigate *multi-player games with goals over finite traces*, which are games where players' goals can be satisfied/achieved after a finite, but arbitrarily large, number of rounds. More specifically, the goals in these games are given by *Linear Dynamic Logic* formulae (LDL_F) which are evaluated over finite sequences of valuations of the variables in Φ , that is, over *finite traces* of valuations, instead of PL formulae (as in BGs) or LTL formulae (as in RMGs and iBGs). Thus, while in game with goals over finite traces a play still is an infinite trace of valuations, the satisfaction of a player's goal may occur after an unbounded but finite number of rounds. This sharply contrasts with the case of goals given by LTL formulae (e.g., as in iBGs and RMGs), where it may be that a player's objective is satisfied only after considering the full infinite trace of valuations. This simple feature has significant implications, since rather complex automata constructions for the analysis of logics and games over infinite traces may become conceptually simpler under this new semantic (logic-based) framework. More importantly, this key observation allows one to define an automata model that exactly characterises the set of Nash equilibria in games with goals given by regular objectives.

There are several reasons to consider LDL_F goals. LDL_F offers great expressive power to our logic-based framework, which is indeed equivalent to monadic second-order logic (MSO). On the other hand, LTL interpreted on finite traces (LTL_F) is as expressive as first-order logic (FOL) over finite traces [?]. This, in turn, implies that, over finite traces, while with LTL_F we can only describe star-free regular languages/properties, with LDL_F we can describe all regular languages/properties—that is, the properties and languages that can be described by regular expressions or finite state automata. Nevertheless, the automata-theoretic approach and complexity results for solving their related decision problems are equivalent, showing that the gain in expressiveness is achieved for free. In this paper, we first define multi-player games with LDL_F goals and then investigate their main game-theoretic properties using a new automata-theoretic approach to reasoning about Nash equilibria. Our technique to reason about equilibria builds on automata constructions originally defined to reason about LDL_F formulae [? ?]. Using this automata-theoretic technique we show a number of subsequent verification and characterisation results, as follows.

¹Although similar, iBGs and RMGs have a number of differences that will be discussed later in the paper.

Firstly, we show that *checking* whether some strategy profile is a Nash equilibrium of a game is a PSPACE-complete problem, thus no harder than LDL_F satisfiability [?]. Secondly, we focus on the NE-NONEMPTINESS problem—which asks for the existence of a Nash equilibrium in a multi-player game succinctly specified by a set of Boolean variables and LDL_F formulae—and show that deciding if a multi-player game with LDL_F goals (whether RMG or iBG) has a Nash equilibrium can be solved in 2EXPTIME, thus no harder than solving LDL_F synthesis [?]. The automata technique we use for this problem also shows that the set of Nash equilibria in these games is ω -regular and can therefore be *characterised* using alternating automata. Thirdly, we also provide complexity results for the main decision problems related to the equilibrium analysis of these games with respect to extensions and restrictions of the initially studied framework. In particular, we show that a small extension of the goal language, which we call *Quantified-Prefix Linear Dynamic Logic* (QPLDL_F), has the same automata-theoretic characteristics as LDL_F , and so it can be studied using the same techniques. Moreover, LDL_F synthesis can be expressed in QPLDL_F , ensuring 2EXPTIME-completeness.

Regarding restrictions on the general framework, we first focus on the problem of reasoning with memoryless strategies. We show, using an automata construction, that the set of Nash equilibria for this games is also ω -regular. However, an alternative procedure for this problem, not based on automata, shows that improved complexity can be obtained when compared with the standard automata techniques to reason about LDL_F . Another restriction on strategies considered in the paper is the one of myopic strategies (which can be used to define all beneficial deviations in a game), in which players perform actions that are independent of the current state of the game execution. We show that games with such a restriction can be solved in EXPSpace. We also consider the much more stable solution concept of *strong Nash equilibrium*, where sets of players in the game are allowed to jointly deviate, and provide an adaptation of the automata-based approach that retains the language characterisation and complexity properties of Nash equilibrium.

A key contribution of this work is that our automata-theoretic approach features two novel properties, within the same reasoning framework. Firstly, it shows that *checking* the existence of Nash equilibria can be reduced to a number of LDL_F synthesis and satisfiability problems—generalising ideas initially used to reason about LTL objectives [?]. Secondly, our automata constructions provide reductions where not only non-emptiness but also language equivalence is preserved. This additionally shows that the set of Nash equilibria in infinite games with regular goals is an ω -regular set, to the best of our knowledge, a semantic *characterisation* not previously known, and which do not immediately follow from other representations of Nash equilibria—see, e.g., [? ? ? ? ?].

Motivation and Previous Work

While studying either multi-player games or LDL_F is interesting in itself, from an AI perspective, our main motivation comes from applications to multi-agent systems. In particular, it has been shown that in many scenarios, for instance in the context of planning AI systems [? ?], while logics like LTL, or even LTL over finite traces (LTL_F), can be used to reason about the behaviour of agents in such AI systems, these logics are not powerful enough to express in a satisfactory way the main features of agents

in such a context. In order to illustrate the use of LDL_F , and motivate even further our work, we will present an example in the next section, where some of the goals either are not expressible in LTL^2 or have a more intuitive specification in LDL_F than in LTL . Together with applications to planning AI systems (see [? ?]), this is an example of another instance where one can see an advantage of using a game with LDL_F goals over a game with LTL goals, instead.

Moreover, regarding previous work, while our model builds on RMGs [?] and iBGs [?], where goals are given by LTL formulae, there are at least two main differences with such work. Firstly, we study scenarios that consider memoryless and myopic strategies, for which results on iBGs have not been investigated³. Secondly, and most importantly, the tools developed in this paper to obtain most of our complexity and characterisation results, are *technically* remarkably different from those used for RMGs and iBGs, specifically, with respect to the techniques used in [? ? ?]. To be more precise, for RMGs and iBGs the main question is reduced to rational synthesis [?], whose solution goes via a parity automaton characterising formulae of an extension of Chatterjee *et al*'s Strategy Logic [?], which leads to an automata construction that can be further optimised if computing Nash equilibria is the only concern. Instead, in our case, we reduce the problem *directly* to a question of automata constructed in a different way. As a consequence, we provide a new set of automata constructions which do not rely on nor relate to those used in rational synthesis, *i.e.*, those used to solve RMGs and iBGs. Our automata constructions are also different from those used by De Giacomo and Vardi in [? ? ?], as described next.

In [? ? ?], De Giacomo and Vardi study the satisfiability and synthesis problems for LDL_F , with and without imperfect information. Because of the (game-theoretic) nature of these two problems, their automata constructions deal with two-player zero-sum turn-based scenarios only. Instead, in our case, we deal with multi-player general-sum concurrent scenarios. This difference leads to a completely different technical treatment/manipulation of the automata that can be initially constructed from LDL_F formulae. In fact, their automata constructions and ours are the same only up to the point where LDL_F formulae are translated into automata—that is, the very first step in a long chain of constructions. Moreover, since De Giacomo and Vardi study synthesis and satisfiability problems (represented by two-player games), whereas we study Nash equilibria (in the context of multi-player games), we are required to have a different technical treatment of the automata involved in the solution of the problems investigated in this paper.

2. Formal Framework

2.1. Linear Dynamic Logic on Finite Traces

In this paper, we consider Linear Dynamic Logic on Finite Traces (LDL_F), a temporal logic introduced in [?] in order to reason about systems whose behaviour

²This fact can be proved from the result that LDL_F is equivalent to Monadic-Second Order logic, while LTL is equivalent to First-Order Logic [?]

³Memoryless and myopic strategies were already studied for RMGs in [?].

can be characterised by sets of finite traces, that is, finite sequences of valuation for the variables of the system.

Definition 1 (Syntax). *The syntax of LDL_F is as follows:*

$$\begin{aligned}\varphi &:= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle \rho \rangle \varphi \mid [\rho] \varphi \\ \rho &:= \psi \mid \varphi? \mid \rho + \rho \mid \rho; \rho \mid \rho^*,\end{aligned}$$

where p is an atomic proposition in Φ ; ψ denotes a propositional formula over the atomic propositions in Φ .

The symbol ρ denotes path expressions, which are regular expressions over propositional formulae ψ , with the addition of the test construct $\varphi?$ from propositional dynamic logic (PDL); and φ stands for LDL_F formulae built by applying Boolean connectives and the modal connectives. Tests are used to insert checks for satisfaction of additional LDL_F formulae.

The classic LTL_F operators can be defined as follows: $\mathbf{X}\varphi \equiv \langle \top \rangle \varphi$; $\mathbf{F}\varphi \equiv \langle \top^* \rangle \varphi$; $\mathbf{G}\varphi \equiv [\top^*] \varphi$; $\varphi_1 \mathbf{U} \varphi_2 \equiv \langle (\varphi_1?)^* \rangle \varphi_2$.

LDL_F formulae are interpreted over finite traces of the form $\pi : \{0, \dots, t\} \rightarrow 2^\Phi$ and an integer $i \in \{0, \dots, t\}$.

Definition 2 (Semantics). *The semantics of LDL_F formulae is as follows:*

- $\pi, i \models \neg\varphi$ if $\pi, i \not\models \varphi$;
- $\pi, i \models \varphi_1 \wedge \varphi_2$ if $\pi, i \models \varphi_1$ and $\pi, i \models \varphi_2$;
- $\pi, i \models \varphi_1 \vee \varphi_2$ if $\pi, i \models \varphi_1$ or $\pi, i \models \varphi_2$;
- $\pi, i \models \langle \alpha \rangle \varphi$ if there exists $j \in \{i, \dots, t\}$ such that $(i, j) \in \mathcal{R}(\alpha, \pi)$ and $\pi, j \models \varphi$;
- $\pi, i \models [\alpha] \varphi$ if for all $j \in \{i, \dots, t\}$, if $(i, j) \in \mathcal{R}(\alpha, \pi)$ then $\pi, j \models \varphi$;

where $\mathcal{R}(\alpha, \pi) \subseteq \mathbb{N} \times \mathbb{N}$ is recursively defined by

- $\mathcal{R}(\psi, \pi) = \{(i, i+1) : \pi(i) \models \psi\}$;
- $\mathcal{R}(\varphi?, \pi) = \{(i, i) : \pi, i \models \varphi\}$;
- $\mathcal{R}(\alpha_1 + \alpha_2, \pi) = \mathcal{R}(\alpha_1, \pi) \cup \mathcal{R}(\alpha_2, \pi)$;
- $\mathcal{R}(\alpha_1; \alpha_2, \pi) = \{(i, j) : \exists k \in \{i, \dots, j\}. (i, k) \in \mathcal{R}(\alpha_1, \pi) \wedge (k, j) \in \mathcal{R}(\alpha_2, \pi)\}$;
- $\mathcal{R}(\alpha^*, \pi) = \{(i, i)\} \cup \{(i, j) : \exists k \in \{i, \dots, j\}. (i, k) \in \mathcal{R}(\alpha, \pi) \wedge (k, j) \in \mathcal{R}(\alpha^*, \pi)\}$.

By $\pi \models \varphi$, we denote the fact that $\pi, 0 \models \varphi$.

2.2. Iterated Boolean Games

We now introduce *iterated Boolean games with goals over finite traces* (iBG_F), which build upon the framework of iBGs [?]. In an iBG_F, players' goals are given by LDL_F formulae interpreted on infinite paths of valuations over a given set of Boolean variables.

An iBG_F is a tuple $\mathcal{G} = \langle \mathbb{N}, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n \rangle$, where $\mathbb{N} = \{1, \dots, n\}$ is a set of players, Φ is a set of Boolean variables, partitioned into n sets Φ_1, \dots, Φ_n , and the goals $\gamma_1, \dots, \gamma_n$ of the game are LDL_F formulae over Φ . In an iBG_F each player i is assumed to control a set of propositional variables Φ_i , in the sense that player i has the power to set the values (true “ \top ” or false “ \perp ”) of each of the variables in Φ_i . An *action* for player i is a possible valuation $v_i \in 2^{\Phi_i}$. An *action vector* $\vec{v} = \langle v_1, \dots, v_n \rangle$ is a collection of actions, one for each player in the game. Every action vector determines an overall valuation for the variables in $\Phi = \bigcup_{i=1}^n \Phi_i$ of the game. An iBG_F is played for an infinite number of rounds, as an iBG, but the goals of the game, which are LDL_F formulae, are interpreted on finite traces of such an infinite run. As a consequence, the satisfaction of a player's goal in the game must occur after a finite, yet arbitrarily large, number of rounds. Due to this, we need to define how an LDL_F formula is satisfied on an iBG_F. The most natural way to do so, also implicitly followed in [?], is to say that an infinite play π^∞ ⁴ satisfies an LDL_F formula φ if and only if there exists $k \in \mathbb{N}$ such that the prefix up to k of π^∞ , denoted by $\pi_{<k}^\infty$, satisfies φ , i.e., $\pi_{<k}^\infty \models \varphi$. We also write $\pi^\infty \models \varphi$ if there is $k \in \mathbb{N}$ such that $\pi_{<k}^\infty \models \varphi$.

Observe that this definition allows a formula and its negation to be satisfied on the same infinite play. Consider, for example, the LDL_F formula $\varphi = \langle \top^* \rangle p$, which is satisfied by all and only finite traces ending with a state labelled by p , and the infinite play $\pi^\infty = (\bar{p}p)^\omega$ which toggles the value of p infinitely often. Clearly, φ is satisfied on every prefix of π^∞ of even length, while $\neg\varphi$ is satisfied on every prefix of π^∞ of odd length. Thus, we obtain that $\pi^\infty \models \varphi$ and $\pi^\infty \models \neg\varphi$. This means that the notion of satisfaction given by $\pi^\infty \models \neg\varphi$ cannot be used in place of $\pi^\infty \not\models \varphi$, as they are not equivalent. We discuss this later in the paper, and show that a small extension of LDL_F, which allows one to quantify over finite plays in a game, can be used to write formulae ψ that equals to the non-satisfaction of φ , i.e., where $\pi^\infty \models \psi$ if and only if $\pi^\infty \not\models \varphi$.

Observe that the set of models of an LDL_F formula φ is of the form $\alpha \cdot (2^\Phi)^\omega$, with α representing the set of finite traces satisfying φ . In [?] it has been proven that α can be represented by a regular expression. This implies, as shown later, that the set of infinite plays satisfying an LDL_F formula can be described in terms of a nondeterministic Büchi word automaton (NBW) built upon the nondeterministic finite word automaton (NFW) accepting α , in which accepting states are constructed so that they are sinks with a self-loop. This makes the expressive power of LDL_F be incomparable with that of LTL when considering infinite words—i.e., infinite plays. Indeed, on the one hand, it is known that LTL cannot express the ω -regular expression $(p \cdot 2^\Phi)^* \cdot (2^\Phi)^\omega$ [?], and, on the other hand, LDL_F cannot express the LTL formula $\mathbf{GF}p$ (“always eventually” p), for which every NBW accepting the set of models cannot be of the form described

⁴In this paper, we denote the infinite plays by π^∞ in order to distinguish them from the finite plays, simply denoted by π .

above. Now, in order to illustrate the concepts introduced so far, and further motivate the iBG_F framework, we present an example where the need for LDL_F goals plays an essential role.

Example 1. Consider a file-sharing network composed by a protocol manager and 2 clients who want to share file₁ of size n_1 packets and file₂ of size n_2 packets, respectively. The clients want to eventually download the other client's file, while the protocol manager wants this transfer of information to happen in a fair way. For instance, the manager wants client 1 to always upload in odd time-steps of the communication, while client 2 to always upload in the even time-steps of the communication protocol. Moreover, the download of a given file can be marked as completed only after the whole number of its packets has been uploaded by the other party.

We can represent this protocol by means of a three agent game with $N = \{0, 1, 2\}$ in which $\Phi_0 = \{d_1, d_2\}$, $\Phi_1 = \{u_1\}$, and $\Phi_2 = \{u_2\}$. Variable u_i being true means that a single packet of file _{i} has been uploaded by agent i , while variable d_i being true means that the download of file _{i} has been completed. Regarding the goals of the agents, we have the following (LDL_F) formulae. The two clients 1 and 2 want to eventually download file₂ and file₁, respectively. Thus, we have $\gamma_1 = \langle \top^* \rangle d_2$ and $\gamma_2 = \langle \top^* \rangle d_1$. Regarding the goal γ_0 of the protocol manager, this has to include several requirements. First of all, it requires that client 1 always uploads in odd time-steps until the download of file₁ has been completed, while client 2 does the same on even time-steps. We can represent these properties with the following LDL_F formulae: $\gamma_{upl_1} = \langle (u_1; \top)^* \rangle d_1$ and $\gamma_{upl_2} = \langle (\top; u_2)^* \rangle d_2$. Note that client 1 has no requirement on the even time-steps, as neither client 2 on odd time-steps. The reader might notice that the properties γ_{upl_1} and γ_{upl_2} can be represented neither in LTL nor in LTL_F , which is the finite trace version of LTL [?]. In addition to this, the protocol manager is in charge of marking the files as completely downloaded at the right time of the execution. This means that variable d_i has to be set to true once the whole amount of packets of file _{i} has been uploaded and not before. We can specify this requirement with the following LDL_F formulae: $\gamma_{com_1} = [((\neg u_1)^*; u_1)^{n_1}] d_1$ and $\gamma_{com_2} = [((\neg u_2)^*; u_2)^{n_2}] d_2$. In order to avoid that the protocol manager wrongly marks the download of a file as completed, we also have the following requirements: $\gamma_{incom_1} = \bigwedge_{n < n_1} \neg \langle (\neg u_1)^*; u_1; (\neg u_1)^* \rangle^n d_1$ and $\gamma_{incom_2} = \bigwedge_{n < n_2} \neg \langle (\neg u_2)^*; u_2; (\neg u_2)^* \rangle^n d_2$. The goal of the protocol manager is therefore given by the conjunction of all these conditions: $\gamma_0 = \gamma_{upl_1} \wedge \gamma_{upl_2} \wedge \gamma_{com_1} \wedge \gamma_{com_2} \wedge \gamma_{incom_1} \wedge \gamma_{incom_2}$. To see that the system we have just described/designed has a stable behaviour, from a game-theoretic point of view, we need the concepts of strategies and Nash equilibria, which are introduced next. Then, we will review this example again later on.

2.3. Simple Reactive Modules Language Games

Simple Reactive Modules [?] is a model specification language that is based on Reactive Modules [?] and has been used to describe multi-player games with LTL goals [? ?]. Reactive Modules games (RMG) are an extension of iBG s in which one can specify constraints on the power that a player has over the variables that such a

player controls⁵. In addition, one can specify multi-player games directly in a high-level description language (which one can then use as the input of a verification tool – Reactive Modules are used, *e.g.*, in MOCHA [?] and PRISM [?]), which is more convenient from a user point of view for modelling purposes.

In an RMG, an agent is mapped to a reactive module, a machinery that dynamically specifies the choices that are available to the associated agent. Formally, a reactive module consists of:

- (i) an *interface*, which defines the module’s name and the set of Boolean variables under the *control* of the module; and
- (ii) a number of *guarded commands*, which define the choices available to the module at every state.

Guarded commands are of two kinds: those used for *initialising* the variables under the module’s control (**init** guarded commands), and those for *updating* these variables subsequently (**update** guarded commands). A guarded command has two parts: a condition part (the “guard”) and an action part, which defines how to update the value of (some of) the variables under the control of a module. The intuitive reading of a guarded command $\varphi \rightarrow a$ is “if the condition φ is satisfied, then *one of the choices available to the module is to execute the action a*”. We note that the truth of the guard φ does not mean that *a will be executed*: only that such a command is *enabled* for execution—it *may be chosen*.

Formally, a guarded command g over the set of Boolean variables Φ is an expression

$$\varphi \rightarrow x_1 := \psi_1; \dots; x_k := \psi_k$$

where φ (the guard) is a propositional logic formula over Φ , each x_i is a controlled variable, and each ψ_i is a propositional logic formula over Φ . Let $guard(g)$ denote the guard of g . Thus, in the above rule, $guard(g) = \varphi$. We require that no variable appears on the left hand side of two assignment statements in the same guarded command. We say that x_1, \dots, x_k are the *controlled variables* of g , and denote this set by $ctr(g)$. If no guarded command of a module is enabled, the values of all variables in $ctr(g)$ are left unchanged; in SRML notation, if needed, `skip` will refer to this particular case.

Formally, an SRML module, m_i , is defined as a triple:

$$m_i = \langle \Phi_i, I_i, U_i \rangle,$$

where:

- $\Phi_i \subseteq \Phi$ is the (finite) set of variables controlled by m_i ;
- I_i is a (finite) set of *initialisation* guarded commands, such that for all $g \in I_i$, we have $ctr(g) \subseteq \Phi_i$; and

⁵Iterated Boolean Games result in the special case of a RMG in which no constraint is specified for every agent.

- U_i is a (finite) set of *update* guarded commands, such that for all $g \in U_i$, we have $\text{ctr}(g) \subseteq \Phi_i$.

Modules can be composed in an intersection manner as follows. For two modules $m_1 = (\Phi_1, I_1, U_1)$ and $m_2 = (\Phi_2, I_2, U_2)$ with $\Phi_1 \cap \Phi_2 = \emptyset$, the product module is $m_1 \otimes m_2 = (\Phi_1 \cup \Phi_2, I_1 \otimes I_2, U_1 \otimes U_2)$ where the \otimes -operator over sets of guards G_1 and G_2 is defined as the set of guards g such that there exist $g_1 \in G_1$ and $g_2 \in G_2$ of the form $\varphi_1 \rightarrow x_1^1 := \psi_1^1; \dots; x_k^1 := \psi_k^1$ and $\varphi_2 \rightarrow x_1^2 := \psi_1^2; \dots; x_k^2 := \psi_k^2$, respectively, such that g is of the form:

$$\varphi_1 \wedge \varphi_2 \rightarrow x_1^1 := \psi_1^1; \dots; x_k^1 := \psi_k^1; x_1^2 := \psi_1^2; \dots; x_k^2 := \psi_k^2$$

An SRML *game* over LDL_F goals (RMG_F) is then defined to be a tuple:

$$\mathcal{G} = \langle N, \Phi, M, \eta, \gamma_1, \dots, \gamma_n \rangle$$

where $N = \{1, \dots, n\}$ is a set of agents, Φ is a set of Boolean variables, $M = \{m_1, \dots, m_M\}$ is a set of modules such that $\Phi_{m_1}, \dots, \Phi_{m_M}$ forms a partition of Φ (so every variable in Φ is controlled by some module, and no variable is controlled by more than one module), and $\eta : M \rightarrow N$ is a function assigning a module to an agent. Finally γ_i is an LDL_F formula associated to agent i .

SRML games can be seen as an extension of iBGs in which the strategic power of agents is not apriori fixed but flexibly allocated according to the evolution of the game. Typically, modules can be used to enforce or prevent agents to adopt a desired/undesired behaviour. As an example, consider the module *toggle* described below:

```

module toggle controls { $p$ }
  init
  []  $\top \rightarrow p := \top$ ;
  []  $\top \rightarrow p := \top$ ;
  update
  []  $p \rightarrow p := \perp$ ;
  []  $\neg p \rightarrow p := \top$ 

```

A player associated to this module is left free to set the value of p on the first round, but then is forced to toggle its value at every round. This kind of constraint on the strategic power of the player associated with agent *toggle* cannot be specified in iBGs.

When an agent i is associated to the set of modules $\eta^{-1}(i)$, then, it can control all the variables that are associated to some module in $\eta^{-1}(i)$. Moreover, at every step of the execution and for every module, it can use one and only one active guarded command. Therefore, every game \mathcal{G} is equivalent to a game \mathcal{G}' where every agent i is associated to the module $\bigotimes_{m_i \in \eta^{-1}(i)} m_i$. Such translation from \mathcal{G} to \mathcal{G}' might produce an exponential blow-up in the representation. However, it is not hard to see that all the techniques used in this paper can be easily applied to a RMG_F game with more than a module associated to an agent, thus avoiding this representation expansion and blow-up in the complexity. For simplicity, we focus on games of this special form and denote them by $\mathcal{G} = \langle N, \Phi, m_1, \dots, m_n, \gamma_1, \dots, \gamma_n \rangle$ where $m_i = \langle \Phi_i, I_i, U_i \rangle$ is the single module associated to agent i , specifying the variable control for the whole set Φ_i .

In this paper, we provide results for both iBGs and RMGs. In particular, we show that, for all the problems addressed here, they have the same computational complexity. Therefore, given that RMGs are a proper extension of iBGs, we show the upper-bound complexity results for RMGs and lower-bound complexity results for iBGs.

2.4. Strategies and Nash Equilibria

Strategies in iBG_F and RMG_F are modelled as *deterministic finite state machines*. Formally a deterministic finite state machine for player i is a tuple $\sigma_i = (S_i, s_i^0, \delta_i, \tau_i)$ where, S_i is a finite set of *internal states*, s_i^0 is the *initial state*, $\delta_i : S_i \times 2^\Phi \rightarrow S_i$ is a *transition function*, and $\tau_i : S_i \rightarrow 2^{\Phi_i}$ is the *action function*. By Str_i we denote the set of possible strategies for player i .

In a RMG_F , a strategy σ_i might not comply with module m_i 's specification. Hence, for this case we need to define a consistency condition between the module and the strategy. We say that σ_i is *compatible* with module m_i if:

1. $\tau_i(s_i^0) = \text{exec}_i(g, \emptyset)$, for some $g \in I_i$; and,
2. for all $v \in 2^\Phi$ and $s, s' \in S_i$, if $s' = \delta_i(s, v)$ then $\tau_i(s') = (\tau_i(s) \setminus \text{ctr}(g)) \cup \text{exec}_i(g, v)$, for some $g \in U_i$ that is enabled by v , *i.e.*, such that $v \models \text{guard}(g)$,

where $\text{exec}_i : (I_i \cup U_i) \times 2^\Phi \rightarrow 2^{\Phi_i}$ is the partial function that determines the value of the Boolean variables at the right-hand side of a guarded command when such a guarded command is enabled by a valuation. Formally, exec_i is defined, for a guarded command $g = \varphi \rightarrow x_1^i := \psi_1^i; \dots; x_k^i := \psi_k^i$ and a valuation v , as $\text{exec}_i(g, v) = \{x_j^i \in \{x_1^i, \dots, x_k^i\} : v \models \psi_j^i\}$. Intuitively, the compatibility ensures the agent to comply with the module's specification. On the one hand, Item 1 states that the guarded command that can be executed on the first state must be an initial command that is enabled on the empty evaluation, that is the evaluation in place before the execution starts. On the other hand, Item 2 deals with the fact that, at every iteration, a strategy must execute a guarded command that is enabled in the current state.

From now on, for the case of RMG_F and when it is clear from the context, we will refer to compatible strategies simply as strategies.

A (total) *strategy profile* is a tuple $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ of strategies, one for each player. We also consider partial strategy profiles. For a given set of players $A \subseteq N$, we use the notation σ_A to denote a tuple of strategies, one for each player in A . Moreover, we use the notation σ_{-A} to denote a tuple of strategies, one for each player in $N \setminus A$. We also use σ_i in place of $\sigma_{\{i\}}$ and $\vec{\sigma}_{-i}$ in place of $\vec{\sigma}_{N \setminus \{i\}}$. Finally, for two strategy profiles $\vec{\sigma}$ and $\vec{\sigma}'$, by $(\vec{\sigma}_A, \vec{\sigma}'_{-A})$ we denote the strategy profile given by associating the strategies in $\vec{\sigma}$ to players in A and strategies in $\vec{\sigma}'$ to players in $N \setminus A$.

Since strategies are deterministic, each profile $\vec{\sigma}$ determines a unique infinite play, denoted by $\pi^\infty(\vec{\sigma})$, which consists of an infinite sequence of valuations, one for each round of the game. Each player i has a preference relation over plays $\pi^\infty \in (2^\Phi)^\omega$, which is determined by its goal γ_i . We say that π^∞ is preferred over $\pi^{\infty'}$ by agent i , and write $\pi^\infty \succeq_i \pi^{\infty'}$, if and only if $\pi^{\infty'} \models \gamma_i$ implies that $\pi^\infty \models \gamma_i$. Using this notion of preference, one can introduce the concept of *Nash Equilibrium*. We say that $\vec{\sigma}$ is a Nash Equilibrium strategy profile if, for each agent i and a strategy $\sigma'_i \in \text{Str}_i$, it holds

that $\pi^\infty(\vec{\sigma}) \succeq_i \pi^\infty(\vec{\sigma}_{-i}, \sigma'_i)$. In addition, by $\text{NE}(\mathcal{G}) \subseteq \text{Str}_1 \times \dots \times \text{Str}_n$ we denote the set of Nash Equilibria of the game \mathcal{G} .

Example 2. Consider again the system in Example 1. A possible strategy σ_1 for player 1 is a finite-state machine that sets variable u_1 to true on odd rounds of the execution, while a strategy σ_2 for player 2 might set u_2 to true on even rounds of the execution. In addition, a possible strategy for player 0, say σ_0 , might be a finite-state machine that sets variable d_i to true only after u_i has been set to true exactly n_i times in the execution. Then, the strategy profile $\vec{\sigma} = (\sigma_0, \sigma_1, \sigma_2)$ will be such that the execution $\pi^\infty = \pi^\infty(\vec{\sigma})$ satisfies γ_0 , γ_1 , and γ_2 , and therefore is a Nash equilibrium. Indeed, checking that a strategy profile is a Nash equilibrium of a game is one of the main concerns of this paper, as formalised next.

Equilibrium Checking. We are interested in a number of questions related to the *equilibrium analysis* of logic-based multi-player games [? ?].

NE MEMBERSHIP. Given a game \mathcal{G} and a strategy profile $\vec{\sigma}$:

Is it the case that $\vec{\sigma} \in \text{NE}(\mathcal{G})$?

which asks if a strategy profile is a Nash equilibrium of a game.

The second decision problem we are interested in is the following:

NE NON-EMPTYNESS. Given a game \mathcal{G} :

Is it the case that $\text{NE}(\mathcal{G}) \neq \emptyset$?

which asks if a given game has at least one Nash equilibrium.

Finally, we also consider two decision problems, that are known in the literature as *equilibrium checking* [?], formally stated as follows:

E/A-NASH. Given a game \mathcal{G} and LDL_F formula φ :

Does $\pi^\infty(\vec{\sigma}) \models \varphi$ hold, for some/all $\vec{\sigma} \in \text{NE}(\mathcal{G})$?

which asks if φ is satisfied by some/every Nash equilibrium of \mathcal{G} .

In the following sections, we study the above questions for both iBG_F and RMG_F , in particular using an automata-theoretic approach.

3. NE Membership

In order to address the NE MEMBERSHIP problem, we first provide some preliminary results on automata. An interested reader can find definitions and more details in [?].

Consider a nondeterministic finite word automaton (NFW) $\mathcal{A} = \langle \Sigma, S, s_0, \varrho, F \rangle$, recognizing a regular language $\mathcal{L}(\mathcal{A})$. Then consider the nondeterministic Büchi word automaton (NBW) $\mathcal{A}^\infty = \langle \Sigma, S, s_0, \varrho', F \rangle$, where, for all σ and s , we have that $\varrho'(\sigma, s) = \varrho(\sigma, s)$, if $s \notin F$, and $\varrho'(\sigma, s) = \{s\}$, otherwise.

Intuitively, the automaton \mathcal{A}^∞ mimics the operations of the automaton \mathcal{A} until an accepting state s is reached. From that point on, the automaton disregards mimicking

\mathcal{A} and starts looping indefinitely over s . Thus, on the one hand, for every finite word $\pi \in \mathcal{L}(\mathcal{A})$, every infinite extension π^∞ , that is, an infinite word such that π is a prefix of it, is accepted by \mathcal{A}^∞ . On the other hand, for an infinite word π^∞ accepted by \mathcal{A}^∞ , there must be a prefix π accepted by \mathcal{A} . This fact can be shown formally with the following theorem.

Theorem 1. *Let \mathcal{A} be a NFW. Then, for all $\pi^\infty \in \Sigma^\omega$, we have that $\pi^\infty \in \mathcal{L}(\mathcal{A}^\infty)$ iff there exists $k \in \mathbb{N}$ such that $\pi = (\pi^\infty)_{\leq k} \in \mathcal{L}(\mathcal{A})$. In particular, we have that $\mathcal{L}(\mathcal{A}_\varphi^\infty) = \{\pi^\infty \in (2^\Phi)^\omega : \pi^\infty \models \varphi\}$*

Proof. The proof proceeds by double implication.

- Assume $\pi^\infty \in \mathcal{L}(\mathcal{A}^\infty)$ and let $\rho^\infty \in (2^S)^\omega$ be an accepting run. Then, since $\varrho'(\sigma, s) = \{s\}$ for every $s \in F$, we have that ρ^∞ is of the form $s_0 \cdot s_1 \cdot \dots \cdot s_{k-1} \cdot s_k^\omega$, with $s_k \in F$ and $s_j \notin F$ for all $j < k$. Then, consider the prefix $\pi = (\pi^\infty)_{\leq k}$. By the definition of ϱ' , it follows that $s_0 \cdot s_1 \cdot \dots \cdot s_{k-1} \cdot s_k$ is an accepting run for π in \mathcal{A} , and so $\pi \in \mathcal{L}(\mathcal{A})$.
- Let $k \in \mathbb{N}$ such that $\pi = (\pi^\infty)_{\leq k} \in \mathcal{L}(\mathcal{A})$ and let $s_0 \cdot s_1 \cdot \dots \cdot s_k$ be an accepting run in \mathcal{A} . Moreover, let $j \in \{0, \dots, k\}$ be such that $s_j \in F$ and $s_h \notin F$ for all $h < j$ ⁶. Then, the infinite sequence $s_0 \cdot s_1 \cdot \dots \cdot s_{j-1} \cdot s_j^\omega$ is an accepting run of π^∞ in \mathcal{A}^∞ . Thus, we have that $\pi^\infty \in \mathcal{L}(\mathcal{A}^\infty)$.

□

To solve the NE MEMBERSHIP problem for RMG_F , we need to account for the fact that strategies adopted by agents must comply with their module specification. In terms of automata, we need to make sure that the accepted language is restricted to the infinite play that can be generated by such complying strategies.

Let $\mathcal{G} = \langle N, \Phi, m_1, \dots, m_n, \gamma_1, \dots, \gamma_n \rangle$ be a RMG_F and consider the automaton $\mathcal{A}_\mathcal{G} = \langle 2^\Phi, S_\mathcal{G}, s_\mathcal{G}^0, F_\mathcal{G}, \varrho_\mathcal{G} \rangle$ defined as:

- $S = 2^\Phi \cup \{\epsilon\}$;
- $s_\mathcal{G}^0 = \epsilon$;
- $F_\mathcal{G} = S_\mathcal{G}$;
- $\varrho_\mathcal{G}(\epsilon, v) = \begin{cases} v, & \text{if } \exists g_1 \in I_1, \dots, g_n \in I_n \emptyset \models \text{guard}(g_i)v = \bigcup_{i \in N} \text{exec}_i(g_i, \emptyset); \\ \emptyset, & \text{otherwise} \end{cases}$;
- $\varrho_\mathcal{G}(s, v) = \begin{cases} v, & \text{if } \exists g_1 \in U_1, \dots, g_n \in U_n \emptyset \models \text{guard}(g_i)v = \bigcup_{i \in N} \text{exec}_i(g_i, \emptyset); \\ \emptyset, & \text{otherwise} \end{cases}$

⁶Such a j exists because s_k is accepting.

Intuitively, at every round of its execution, the automaton stores the current evaluation of the game in its state s . Then, when an evaluation v is sent to it, the automaton moves to the state v itself if, and only if, there exists a tuple of commands g_1, \dots, g_n that are enabled in s and whose combined executions leads the game from s to v . For the case no sets of enabled commands can lead to v , the automaton rejects the path, as it has found an illegal move in it. More formally, we have the following result.

Lemma 1. *Let \mathcal{G} be a $\text{RMG}_{\mathbb{F}}$ and $\mathcal{A}_{\mathcal{G}}$ its corresponding automaton. Then, it holds that $\mathcal{L}(\mathcal{A}_{\mathcal{G}})$ is exactly the set of possible executions in \mathcal{G} .*

Proof. We prove the lemma by double inclusion. First, assume that π^∞ is a play in \mathcal{G} and show that it is accepted by $\mathcal{A}_{\mathcal{G}}$. Let s_0, s_1, \dots be the run of $\mathcal{A}_{\mathcal{G}}$ over π^∞ . By induction on h , we prove that $s_{h+1} = \pi_h^\infty$, for every $h \in \mathbb{N}$ and so that the run is accepting. As base case, for $h = 0$, we have that $s_1 = \varrho_{\mathcal{G}}(s_0, \pi_0^\infty) = \varrho_{\mathcal{G}}(\epsilon, \pi_0^\infty)$ and, since π^∞ is a legal execution in \mathcal{G} , there exist $g_1 \in I_1, \dots, g_n \in I_n$, with $\emptyset \models \text{guard}(g_i)$, for all $i \in N$, such that $\pi_1^\infty = \bigcup_{i \in N} \text{exec}_i(g_i, \emptyset)$, this implying $\varrho_{\mathcal{G}}(\epsilon, \pi_0^\infty) = \pi_0^\infty$, proving the statement for the base case. For the induction case, let $h \geq 0$ and assume $s_{h+1} = \pi_h^\infty$. We have to prove that $s_{h+2} = \pi_{h+1}^\infty$. We have that $s_{h+2} = \varrho_{\mathcal{G}}(s_{h+1}, \pi_{h+1}^\infty)$. Moreover, since π^∞ is an execution in \mathcal{G} , there exist $g_1 \in U_1, \dots, g_n \in U_n$, with $\pi_h^\infty \models \text{guard}(g_i)$, for all $i \in N$, such that $\pi_{h+1}^\infty = \bigcup_{i \in N} \text{exec}_i(g_i, \emptyset)$. Now, observe that by induction hypothesis, we have that $s_{h+1} = \pi_h^\infty$, and so it follows that $s_{h+2} = \varrho_{\mathcal{G}}(s_{h+1}, \pi_{h+1}^\infty) = \varrho_{\mathcal{G}}(\pi_h^\infty, \pi_{h+1}^\infty) = \pi_{h+1}^\infty$, the last equality following from the definition of $\varrho_{\mathcal{G}}$.

For the other direction, let us assume $\pi^\infty \in \mathcal{L}(\mathcal{A}_{\mathcal{G}})$ and prove that π^∞ is an execution in \mathcal{G} . We have to show prove that there exist $g_1 \in I_1, \dots, g_n \in I_n$ such that $\emptyset \models \text{guard}(g_i)$, for all $i \in N$ and $\pi_0^\infty = \bigcup_{i \in N} \text{exec}_i(g_i, \emptyset)$ and that, for all $h \in \mathbb{N}$, there exist $g_1^h \in U_1, \dots, g_n^h \in U_n$ such that $\pi_h^\infty \models \text{guard}(g_i^h)$, for all $i \in N$ and $\pi_{h+1}^\infty = \bigcup_{i \in N} \text{exec}_i(g_i^h, \pi_h^\infty)$. But this clearly follows from the definition of $\varrho_{\mathcal{G}}$ and the fact that the run of $\mathcal{A}_{\mathcal{G}}$ over π^∞ is accepting. \square

We can now address NE MEMBERSHIP. We show that this problem is PSPACE-complete; for the membership argument, we employ an automata-based algorithm for checking membership. We first introduce, for a given (machine) strategy $\sigma_i = \langle S_i, s_i^0, \delta_i, \tau_i \rangle$ for a player i , a corresponding DFW $\mathcal{A}(\sigma_i) = \langle \Sigma, Q_i, q_i^0, \varrho_i, F_i \rangle$ where: $\Sigma = 2^\Phi$ is the alphabet set, $Q_i = (S_i \times 2^\Phi) \cup \{\text{sink}\}$ is the state set, where $\text{sink} \notin S_i \times 2^\Phi$ is a fresh state, $q_i^0 = (s_i^0, \emptyset)$ is the initial state, $F_i = S_i \times 2^\Phi$ is the final state set, and ϱ_i is the transition relation such that, for all $(s, v) \in S_i \times 2^\Phi$ and $v' \in \Sigma$,

- $\varrho_i((s, v), v') = \begin{cases} \delta_i(s, v), & \text{if } \tau_i(s) = v' \upharpoonright_{\Phi_i} \\ \text{sink}, & \text{otherwise} \end{cases}$, and
- $\varrho_i(\text{sink}, v') = \text{sink}$

Let $\mathcal{L}(\mathcal{A}(\sigma_i))$ denote the set of infinite words in $(2^\Phi)^\omega$ accepted by $\mathcal{A}(\sigma_i)$. It is easy to see that such a set is exactly the same set of plays that are possible outcomes in a game where player i uses strategy σ_i . Similarly, for a given set of players $A \subseteq N$ and a partial strategy profile $\vec{\sigma}_A$, we have that, for $\mathcal{A}(\vec{\sigma}_A) = \bigotimes_{i \in A} \mathcal{A}(\sigma_i)$, the product of these automata, the language $\mathcal{L}(\mathcal{A}(\vec{\sigma}_A))$ contains exactly those infinite plays in a game where players in A play according to the strategies given in $\vec{\sigma}_A$. Moreover,

in [?] it is shown, for every LDL_F formula φ , how to build and check on-the-fly a NFW $\mathcal{A}_\varphi = \langle S, 2^\Phi, \{s_0\}, \delta, \{s_f\} \rangle$, such that, for every finite trace $\pi \in (2^\Phi)^*$, we have $\pi \models \varphi$ if and only if $\pi \in \mathcal{L}(\mathcal{A}_\varphi)$, where by $\mathcal{L}(\mathcal{A}_\varphi)$ we denote the language of finite words (that is, the language of finite traces over 2^Φ) accepted by the automaton \mathcal{A}_φ . Such a construction makes use of a function δ simulating the transition relation of the corresponding alternating finite word automaton (AFW), which takes a subformula ψ of φ and a valuation of variables $\Pi \subseteq \Phi$, and recursively returns a combination of subformulae. A suitable modification of such an algorithm allows one to construct the NBW $\mathcal{A}_\varphi^\infty$. As a matter of fact, observe that the only final state s_f of the automaton \mathcal{A}_φ built in [?] does not have any outgoing transition. Then, given the construction of $\mathcal{A}_\varphi^\infty$, we only need to add a loop to it, for every possible valuation.

Input: an LDL_F formula φ and an NBW $\mathcal{A} = \langle 2^\Phi, Q, q_0, \varrho, F \rangle$.
Output: NBW $\mathcal{A}_\varphi^\infty \times \mathcal{A} = \langle 2^\Phi, S', \{s'_0\}, F', \varrho' \rangle$.
 $s'_0 \leftarrow \{(\varphi, q_0, 1)\}$ $F' \leftarrow \{\emptyset\} \times Q \times \{1\}$
 $S \leftarrow \{s_0\} \cup F'$ $\varrho \leftarrow \{((\emptyset, q, 1), \Pi, (\emptyset, q', 2)) : \Pi \in 2^\Phi \wedge q' \in \varrho(q, \Pi)\} \cup$
 $\{((\emptyset, q, 2), \Pi, (\emptyset, q', 2)) : \Pi \in 2^\Phi \wedge q' \in \varrho(q, \Pi) \wedge q \in Q \setminus F'\} \cup \{((\emptyset, q, 2), \Pi, (\emptyset, q', 1))$
 $: \Pi \in 2^\Phi \wedge q' \in \varrho(q, \Pi) \wedge q \in F'\}$
while (S' or ϱ' change) **do**
 for $s \in S'$, $q \in Q$ and $\Pi \in 2^\Phi$ **do**
 for $q' \subseteq CL(\varphi)$ and $q' \in \varrho(q, \Pi)$ **do**
 if $s' \models \bigwedge_{\psi \in q} \delta(\psi, \Pi)$ **then**
 if $q \in F'$ **then**
 $S' \leftarrow S' \cup \{(s, q, 2), (s', q', 1)\}$ $\varrho' \leftarrow \varrho' \cup \{((s, q, 2), \Pi, (s', q', 1))\}$
 else
 $S' \leftarrow S' \cup \{(s', q', 1), (s', q', 2), (s, q, 2)\}$ $\varrho' \leftarrow \varrho' \cup$
 $\{((s, q, 2), \Pi, (s', q', 1))\}$

Algorithm 1: Intersection construction.

However, it cannot be used as it is to obtain the PSPACE complexity for the NE MEMBERSHIP problem. Indeed, we need to combine the NBW $\mathcal{A}_\varphi^\infty$ with the automata $\mathcal{A}_{\vec{\sigma}}$ and $\mathcal{A}_{\vec{\sigma}_{-i}}$ provided by the NE MEMBERSHIP problem instance. To do this, we need to adapt the construction in order to handle these products. Note that both $\mathcal{A}_{\vec{\sigma}}$ and $\mathcal{A}_{\vec{\sigma}_{-i}}$ can be considered as NBW. Thus, it is enough to deliver an algorithm that builds an automaton intersection between $\mathcal{A}_\varphi^\infty$ and a generic NBW \mathcal{A} .

<p>Input: a RMG_F \mathcal{G} and a strategy profile $\vec{\sigma}$.</p> <p>Output: “Yes” if $\vec{\sigma} \in \text{NE}(\mathcal{G})$; “No” otherwise.</p> <p>if $\mathcal{L}(\mathcal{A}(\vec{\sigma})) \not\subseteq \mathcal{L}(\mathcal{A}_{\mathcal{G}})$ then</p> <p> \perp return “Error”</p> <p>for $i \in N$ do</p> <p> if $\mathcal{L}(\mathcal{A}(\vec{\sigma}) \otimes \mathcal{A}_{\gamma_i}^\infty) = \emptyset$ then</p> <p> if $\mathcal{L}(\mathcal{A}(\vec{\sigma}_{-i}) \otimes \mathcal{A}_{\mathcal{G}} \otimes \mathcal{A}_{\gamma_i}^\infty) \neq \emptyset$ then</p> <p> \perp return “No”</p> <p>return “Yes”</p>

Algorithm 2: NE Membership for RMG_F .

When one of the two NBW involved in the product derives from an LDL_F formula φ , we can adapt the on-the-fly construction provided in [?], as described in Algorithm 1.

With this construction in place, one can show that Algorithm 2 runs in PSPACE and solves NE MEMBERSHIP for RMG_F . In particular, the algorithm checks, using the automata constructions presented before, whether a strategy profile is a Nash equilibrium by checking for beneficial deviations in the game for every player.

Theorem 2. *The NE MEMBERSHIP problems for $i\text{BG}_F$ and RMG_F are PSPACE-complete.*

Proof. To show that Algorithm 2 is correct, assume that the algorithm returns “Yes” on a given instance $(\mathcal{G}, \vec{\sigma})$. This means that it does not return the “Error” message in the initial conditional check. This means that the outcome of $\vec{\sigma}$ is in the language of $\mathcal{A}_{\mathcal{G}}$ and so the strategies are complying with their modules specifications. Moreover, the algorithm does not return “NO”, which means that, for every agent i , either the innermost or the outermost conditional checks are false. In case the outermost is false, then we have that $\mathcal{L}(\mathcal{A}(\vec{\sigma})) \cap \mathcal{L}(\mathcal{A}_{\gamma_i}) \neq \emptyset$, meaning that the play $\pi^\infty(\vec{\sigma})$ is such that $\pi^\infty(\vec{\sigma}) \models \gamma_i$. Thus, player i is satisfied in the context $\vec{\sigma}$ and so it does not have any incentive to deviate from it. On the other hand, if the outermost returns true but the innermost returns false, then we have that $\mathcal{L}(\mathcal{A}(\vec{\sigma}_{-i}) \otimes \mathcal{A}_{\gamma_i}) \cap \mathcal{L}(\mathcal{A}_{\mathcal{G}}) = \emptyset$, which means that the satisfaction of γ_i is incompatible with the partial strategy profile $\vec{\sigma}_{-i}$, no matter how player i behaves compatibly with its modules specification. This, in terms of strategies, implies that there is no beneficial deviation for player i to get its goal achieved. Hence, the strategy profile $\vec{\sigma}$ is a Nash equilibrium of the game.

On the other hand, assume $\vec{\sigma}$ is a Nash equilibrium. Then, no player i has an incentive to deviate. This can be the case for two reasons: either $\pi^\infty(\vec{\sigma}) \models \gamma_i$, or there is no compatible strategy σ'_i such that $\pi^\infty(\vec{\sigma}_{-i}, \sigma'_i) \models \gamma_i$. If the former, then we have that $\mathcal{L}(\mathcal{A}(\vec{\sigma})) \cap \mathcal{L}(\mathcal{A}_{\gamma_i}) \neq \emptyset$ and so the check on the outermost conditional is false. If the latter, then it follows that $\mathcal{L}(\mathcal{A}(\vec{\sigma}_{-i}) \otimes \mathcal{A}_{\gamma_i}) \cap \mathcal{L}(\mathcal{A}_{\mathcal{G}}) = \emptyset$, making the check on innermost conditional is false. Since this reasoning holds for every player i , it can be concluded that Algorithm 2 ends by returning “Yes”, which concludes the proof of correctness.

Regarding the complexity, note that all the conditional checks involve a nonemptiness test of NFW built by means of the Algorithm 1, whose complexity is PSPACE.

Since this procedure is called n times, where n is the number of agents, we obtain a PSPACE upper bound.

We now show hardness on iBG_F by providing a reduction from the satisfiability problem of LDL_F formulae, which is known to be PSPACE-complete [?]. Consider an LDL_F formula φ and then define the one-player game \mathcal{G} , with player set $\{1\}$, in which player 1 controls all the variables in φ plus an additional variable $\{p\}$ that does not appear in φ , and whose goal is $\gamma_1 = \varphi \wedge p$. Moreover, let σ be the strategy for player 1 that myopically plays \emptyset in all rounds. Then, such a strategy, which clearly is linear in the size of φ since it has constant size, is such that $\sigma \models \neg\gamma_1$. Now, we will show based on this reduction that φ is satisfiable if and only if $\sigma \notin \text{NE}(\mathcal{G})$. Firstly, if φ is satisfiable then player 1 can (beneficially) deviate to a myopic strategy, say σ' , that generates an infinite play with p in the first round and such that one of its prefixes satisfies φ —in which case $\sigma' \models \gamma_1$. Then, $\sigma \notin \text{NE}(\mathcal{G})$. On the other hand, if φ is not satisfiable, then γ_1 is not satisfiable either. Then, it is clear that there is no strategy σ' to which player 1 can beneficially deviate to achieve its goal; hence σ is a Nash equilibrium of \mathcal{G} . Because PSPACE is closed under complement, PSPACE-hardness follows. \square

4. NE Non-Emptiness and Equilibrium Checking Problems

Now, let us study NE NON-EMPTYNESS for both iBG_F and RMG_F . We first prove a result for iBG_F and then how to adapt it for the case of RMG_F . Also in this case we use an automata-theoretic approach. We show how, given a game \mathcal{G} , it is possible to construct an alternating automaton $\mathcal{A}_{\text{NE}}(\mathcal{G})$ such that $\mathcal{A}_{\text{NE}}(\mathcal{G})$ accepts precisely the set of plays that are generated by the Nash equilibria of \mathcal{G} . A distinguishing feature of our automata technique is that it is *language preserving*, that is, $\mathcal{A}_{\text{NE}}(\mathcal{G})$ recognizes exactly the set of plays that are obtained by some Nash equilibrium in the game. Hereafter, we call *Nash runs* the elements in such a set of runs. This property of our construction is the key to show that the set of Nash runs is, in fact, ω -regular. Also, note that as we now have to find (and not simply check) a strategy profile, we cannot use the automata of the form $\mathcal{A}(\sigma_i)$ provided above, as there is no known strategy σ_i , for each player i , that can be used here.

First, we recall the characterisation of Nash equilibria provided in [?]. For a given $RMG_F \mathcal{G} = \langle N, \Phi, m_1, \dots, m_n, \gamma_1, \dots, \gamma_n \rangle$ and a designated player $j \in N$, we say that $\vec{\sigma}_{-j}$ is a punishment profile against j if, for every strategy σ'_j , it holds that $(\vec{\sigma}_{-j}, \sigma'_j) \not\models \gamma_j$. In [?], it has been proven that $\vec{\sigma} \in \text{NE}(\mathcal{G})$ if and only if there exists $W \subseteq N$ such that $\sigma \models \gamma_i$ for every $i \in W$ and, for every $j \in L = N \setminus W$, the profile $\vec{\sigma}_{-j}$ is a punishment strategy against j , that is, a winning strategy profile of the coalition of players N_{-j} for the negation of the goal of player j .

Thus, we can think of finding punishment strategies in terms of synthesizing a finite state machine controlling Φ_{-j} . To do this, we apply an automata-theoretic approach. First of all, we build the alternating Rabin word automaton (ARW) \mathcal{A}_{γ_j} , used to recognize the models of γ_j , and then the product $\mathcal{A}_{\gamma_j}^{\mathcal{G}} = \mathcal{A}_{\gamma_j} \oplus \mathcal{A}_{\mathcal{G}}$, filtering the models that are compatible with an execution of the $RMG_F \mathcal{G}$. Analogously, the automaton $\mathcal{A}_{\gamma_i}^{\mathcal{G}} = \overline{\mathcal{A}_{\gamma_i}} \otimes \mathcal{A}_{\mathcal{G}}$ recognizes the plays that both are compatible with the $RMG_F \mathcal{G}$ and do not satisfy γ_i . At this point, by means of Theorem 2 in [?], we build a nondeterministic

Rabin automaton on trees (NRT) $\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}}$ that recognizes exactly those trees T that are obtained from an execution of a compatible winning strategy of the coalition N_{-j} when the goal is to avoid the satisfaction of γ_j . Now, following Corollary 17 in [?], we can build a NRW $\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}''}}$ such that $\mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}''}}) = \{\pi^\infty \in (2^\Phi)^\omega : \exists T \in \mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}}), \pi^\infty \subseteq T\}$, where by $\pi^\infty \subseteq T$ we denote the fact that π^∞ is a branch of the tree T starting at its root.

Now, let us fix $W \subseteq N$ for a moment, and consider the product automaton $\mathcal{A}_{\overline{L}} = \bigotimes_{j \in \overline{L}} \overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}''}}$. By the semantics of the product operation we obtain that $\mathcal{A}_{\overline{L}}^{\mathcal{G}}$ accepts those paths that are generated by some punishment profile, compatible with \mathcal{G} , for each $j \in \overline{L}$. Moreover, consider the automaton $\mathcal{A}_W = \bigotimes_{i \in W} \mathcal{A}_{\gamma_i}^{\mathcal{G}}$, recognizing the paths that satisfy every γ_i , for $i \in W$. Thus, we have that the product automaton $\mathcal{A}_W \otimes \mathcal{A}_{\overline{L}}$ accepts exactly those paths for which every γ_i , with $i \in W$, is satisfied while, for each $j \in \overline{L}$ the coalition N_{-j} is using a punishment strategy against j . Now, in order to exploit the characterisation given in [?], we only need to quantify over $W \subseteq N$. This, in terms of automata, corresponds to the union operation. Then, we get the following automata characterisation:

$$\mathcal{A}_{\text{NE}}(\mathcal{G}) = \bigoplus_{W \subseteq N} (\mathcal{A}_W \otimes \mathcal{A}_{\overline{L}}).$$

Theorem 3 (RMG_F Expressiveness). *For a RMG_F game \mathcal{G} , the automaton $\mathcal{A}_{\text{NE}}(\mathcal{G})$ recognizes the set of Nash runs of \mathcal{G} . Therefore, the set of Nash runs of \mathcal{G} is ω -regular.*

Proof. We prove the theorem by double implication. From left to right, assume that $\pi^\infty \in \mathcal{L}(\mathcal{A}_{\text{NE}}(\mathcal{G}))$. Then, there is $W \subseteq N$ such that $\pi^\infty \in \mathcal{L}(\mathcal{A}_W \otimes \mathcal{A}_{\overline{L}})$. Observe that, w.l.o.g. we can assume that π^∞ is an ultimately periodic play [?] and so that there exists a finite-state machine $\Delta_{\pi^\infty} = (Q_{\pi^\infty}, q_{\pi^\infty}^0, \delta_{\pi^\infty}, \tau_{\pi^\infty})$, controlling all the variables in Φ , i.e., $\tau_{\pi^\infty} : Q_{\pi^\infty} \rightarrow 2^\Phi$, that generates π^∞ . Moreover, observe that, for each $j \in \overline{L}$, $\pi^\infty \in \mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}''}})$ implies that there exists $T_j \in \mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}})$ such that $\pi^\infty \subseteq T_j$. This implies that, for each $j \in \overline{L}$, there is a finite-state machine $\Delta_j = (Q_j, q_j^0, \delta_j, \tau_j)$, controlling all the variables but Φ_j , i.e., $\tau_j : Q_j \rightarrow 2^{\Phi - \Phi_j}$, that generates the branches of T_j , according to the output of variables in Φ_j , including π^∞ . Now, for each $i \in N$, define the strategy $\sigma_i = (S_i, s_i^0, \delta_i, \tau_i)$ as follows:

- $S_i = Q_{\pi^\infty} \times \prod_{j \in \overline{L}} Q_j \times (L \cup \{\top\})$ is the product of the state-space of Δ_{π^∞} together with the state-space of each Δ_j , for each $j \in \overline{L}$, plus a flag component given by $L \cup \{\top\}$;
- $s_i^0 = (q_{\pi^\infty}^0, q_{j_1}^0, \dots, q_{j_{|\overline{L}|}}^0, \top)$, collecting all the initial states of the finite state machines, Δ_{π^∞} and Δ_j , for each $j \in \overline{L}$, flagged with the symbol \top ;
- δ_i is defined as follows: for each $(q, q_{j_1}, \dots, q_{j_{|\overline{L}|}}, \top)$ and $v \in 2^\Phi$, $\delta_i((q, q_{j_1}, \dots, q_{j_{|\overline{L}|}}, \top), v) = (\delta_{\pi^\infty}(q, v), \delta_{j_1}(q_{j_1}, v), \dots, \delta_{j_{|\overline{L}|}}(q_{j_{|\overline{L}|}}, v), \text{flag})$, where $\text{flat} = \top$ if $v = \tau_{\pi^\infty}(q)$ and $\text{flat} = j$ if $v_{-j} = (\tau_{\pi^\infty}(q))_{-j}$ and $v_j \neq (\tau_{\pi^\infty}(q))_j$.
- $\tau_i((q, q_{j_1}, \dots, q_{j_{|\overline{L}|}}, \top)) = (\tau_{\pi^\infty}(q))_i$ and $\tau_i((q, q_{j_1}, \dots, q_{j_{|\overline{L}|}}, j)) = (\tau_j(q))_i$, for each $j \in \overline{L}$.

Intuitively, a strategy σ_i for player i runs in parallel the i -th component of the finite-state machine Δ_{π^∞} together with the i -th components of the finite-state machines Δ_j that win against the deviating players in L . Note that, by construction, as long as nobody deviates, the outcome of every single Δ_j corresponds to the one of Δ_{π^∞} . We have that the strategy profile $\vec{\sigma}$, given by the union of the strategies defined above, generates π^∞ , and, as soon as a unilateral deviation occurs from player $j \in L$, the partial strategy profile $\vec{\sigma}_{-j}$ starts following the finite-state machine Δ_j , which is by definition winning against j . Thus, $\vec{\sigma}$ is a Nash equilibrium.

From right to left, assume that π^∞ is a Nash run and let $\vec{\sigma}$ be a Nash equilibrium such that $\pi^\infty(\vec{\sigma}) = \pi^\infty$. Moreover, let $W = \{i \in N : \pi^\infty \models \gamma_i\}$. We show that $\pi^\infty \in \mathcal{L}(\mathcal{A}_W \otimes \mathcal{A}_{\bar{L}})$. Since $\pi^\infty \models \gamma_i$, for each $i \in W$, we have that $\pi^\infty \in \mathcal{L}(\mathcal{A}_W)$. Moreover, let $j \in L$. It holds that j does not have a beneficial deviation from $\vec{\sigma}$ and so we have that $\vec{\sigma}_{-j}$ is a winning strategy against j . From the definition of $\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}}$ we have that the tree-execution T_j generated by $\vec{\sigma}_{-j}$ is in $\mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}})$. Now, since $\pi^\infty \subseteq T_{-j}$, we have that $\pi^\infty \in \mathcal{L}(\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}''}})$, for each $j \in L$, implying that $\pi^\infty \in \mathcal{L}(\mathcal{A}_{\bar{L}})$. Hence, we have that $\pi^\infty \in \mathcal{L}(\mathcal{A}_W) \cap \mathcal{L}(\mathcal{A}_{\bar{L}}) = \mathcal{L}(\mathcal{A}_W \otimes \mathcal{A}_{\bar{L}})$, as required. \square

Using Theorem 3 we can address the problem of deciding if a game admits a Nash equilibrium by checking $\mathcal{A}_{\text{NE}}(\mathcal{G})$ for emptiness. Regarding the complexity of building $\mathcal{A}_{\text{NE}}(\mathcal{G})$, observe that the construction of each automaton $\overline{\mathcal{A}_{\gamma_j}^{\mathcal{G}'}}$, provided in [?], is of size doubly exponential with respect to $|\gamma_j|$. Moreover, all the other operations used to build $\mathcal{A}_{\text{NE}}(\mathcal{G})$ involve union and intersection of Rabin automata, which can be performed in time polynomial in the size of the constituting components. This shows that $\mathcal{A}_{\text{NE}}(\mathcal{G})$ is a nondeterministic Rabin automaton on words of size doubly exponential with respect to the game \mathcal{G} . Since checking emptiness of a NRW can be done in NLOGSPACE, we obtain the following result.

Theorem 4. *NE NON-EMPTINESS of RMG_F can be solved in 2EXPTIME.*

Now, to show that E-NASH and A-NASH are in 2EXPTIME, we can also apply an automata-theoretic approach. Indeed, for the E-NASH case, consider a game \mathcal{G} and an LDL_F formula φ . Then, the automaton $\mathcal{A}_\varphi \otimes \mathcal{A}_{\text{NE}}(\mathcal{G})$ recognizes all the plays that both satisfy φ and are a Nash run. Thus, checking the E-NASH problem corresponds to checking the nonemptiness of such automaton. On the other hand, for the A-NASH problem, consider the automaton $\overline{\mathcal{A}_\varphi} \otimes \mathcal{A}_{\text{NE}}(\mathcal{G})$. This product automaton recognizes all plays that do not satisfy the formula φ and are a Nash run. Thus, checking the A-NASH problem corresponds to checking the emptiness of such an automaton. The two constructions above show that both E-NASH and A-NASH can be solved in 2EXPTIME. Formally, combining the results above, we also obtain the following theorem:

Theorem 5. *E-NASH and A-NASH for RMG_F can be solved in 2EXPTIME.*

5. Extensions and Restrictions

We now investigate on some extensions and restrictions on the problems studied in the previous section. As a first result, we show that an extension of the LDL_F language

used to represent players' goals can be used to encode LDL_F synthesis, studied in [?], as a NE NON-EMPTINESS problem. Subsequently, we restrict to two classes of strategies, namely *memoryless* and *myopic* strategies. With respect to memoryless strategies, we show that our automata-based techniques can be used to show that the set of Nash runs for games of this kind is also ω -regular, as in the original problem. An EXPSPACE brute-force approach can be used to show that the induced automata are suboptimal from a complexity point of view.⁷ However, the construction is still based on a simple extension of automata on finite words, making it potentially useful in practice. The case of myopic strategies, instead, is studied using a reduction to the satisfiability problem for the 1-alternation fragment of QPTL, known to be solvable in EXPSPACE [?].

Games with Quantified prefix LDL_F Goals. The results obtained so far show that checking whether a game has a Nash equilibrium can be solved in 2EXPTIME. We now show that an extension of the logic LDL_F , which we call *quantified prefix LDL_F* (QPLDL_F) can also be solved using the same automata-theoretic technique, with the same complexity, and can be used to represent the LDL_F synthesis problem, which is 2EXPTIME-complete. Then, NE NON-EMPTINESS with respect to such an extension is 2EXPTIME-complete.

Syntactically, a QPLDL_F formula φ is obtained from an LDL_F formula ψ by simply adding either an existential \exists or a universal \forall quantifier in front of it, i.e., $\varphi = \exists\psi$ or $\varphi = \forall\psi$. Such a quantification ranges over the set of prefixes of a given infinite path of valuations. Formally, we have that, for a given QPLDL_F formula of the form $\exists\psi$ and an infinite path π^∞ , we have that $\pi^\infty \models \exists\psi$ if there is $k \in \mathbb{N}$ such that $\pi_{<k}^\infty \models \psi$. Analogously, for a QPLDL_F formula of the form $\forall\psi$, we have $\pi^\infty \models \forall\psi$ if $\pi_{<k}^\infty \models \psi$, for all $k \in \mathbb{N}$.

The reader might note that $\exists\psi$ is equivalent to ψ on infinite plays. This means that the set of models for $\exists\psi$ corresponds to the set of infinite models of ψ and so the automaton $\mathcal{A}_{\exists\psi} = \mathcal{A}_\psi$ recognizes the models of $\exists\psi$. Moreover, observe that, for every LDL_F formula ψ and an infinite play π^∞ , we have that $\pi^\infty \models \forall\psi$ iff $\pi^\infty \not\models \exists\neg\psi$. This means that, in order to build the automaton $\mathcal{A}_{\forall\psi}$ for a formula of the form $\forall\psi$, one can first consider the formula $\neg\psi$ and build the corresponding automaton $\mathcal{A}_{\exists\neg\psi}$. It follows that $\mathcal{L}(\mathcal{A}_{\exists\neg\psi})$ is the set of infinite plays that satisfy $\exists\neg\psi$, which is the complement of the set of plays satisfying $\forall\psi$. Thus, $\mathcal{A}_{\forall\psi} = \overline{\mathcal{A}_{\exists\neg\psi}}$. Using these constructions one can solve NE NON-EMPTINESS, E-NASH, and A-NASH with QPLDL_F goals by applying the same automata-theoretic technique used for LDL_F . Then, we have the following result.

Theorem 6. *NE NON-EMPTINESS, E-NASH, and A-NASH with QPLDL_F goals, for both iBG and RMG can be solved in 2EXPTIME. Moreover, the sets of Nash equilibria for these classes of games is ω -regular.*

To obtain a matching lower bound, observe that, given the interpretation of QPLDL_F formulae, it is possible to encode the synthesis problem for LDL_F formulae as presented

⁷Of course, with respect to EXPRESSIVENESS this is an irrelevant feature of the automata construction.

in [?]. Indeed, in such a case we only have to set a two-player game \mathcal{G} in which, say Player 1, controls the same variable as the system for the synthesis problem, and Player 2 controls the environment variables. At this point, by setting $\gamma_1 = \exists\psi$ and $\gamma_2 = \forall\neg\psi$, one ensures that Player 1 and Player 2 have exactly the same behaviours of system and environment in the synthesis problem, respectively. In addition to Player 1 and Player 2, to ensure a reduction to NE NON-EMPTINESS one can add two players that trigger a “matching pennies” game in case ψ is not synthesised. With this reduction it follows that NE NON-EMPTINESS is 2EXPTIME-complete.

Formally, consider an LDL_F formula φ and the synthesis problem for it, in which the system controls a set of (output) variables X while the environment controls a set of (input) variables Y . Then, consider the four-player $\text{iBG}_F \mathcal{G}_\varphi$ with QPLDL_F goals such that:

- Player 1 controls X and has $\gamma_1 = \exists\varphi$ as goal;
- Player 2 controls Y and has $\gamma_2 = \forall\neg\varphi$ as goal;
- Player 3 controls a fresh Boolean variable p and has $\gamma_3 = \exists\varphi \vee (p \leftrightarrow q)$ as goal; and
- Player 4 controls a fresh Boolean variable q and has $\gamma_4 = \exists\varphi \vee \neg(p \leftrightarrow q)$ as goal.

Using the above construction, we can show that the synthesis problem for an LDL_F formula φ can be solved by addressing the NE NON-EMPTINESS problem for \mathcal{G}_φ , from which we derive the following theorem.

Theorem 7. *NE NON-EMPTINESS, E-NASH, and A-NASH are 2EXPTIME-complete for both iBG_F and RMG_F .*

In fact, Theorem 7 is proved using the lemma given below.

Lemma 2. *The synthesis problem for an LDL_F formula φ over a set of Boolean variables $X \cup Y$, where the system controls the variables in X and the environment the variables in Y has a positive answer if and only if the game \mathcal{G}_φ has a Nash equilibrium.*

Proof. We prove the lemma for iBG_F by double implication. From left to right, assume that the system has a winning strategy σ_1 against the environment in the synthesis problem. Then every strategy profile $\vec{\sigma}$ in \mathcal{G}_φ in which Player 1 uses σ_1 is a Nash equilibrium. Indeed, as σ_1 is a winning strategy for the synthesis problem, we have that $\vec{\sigma} \models \varphi$ and so $\vec{\sigma} \models \exists\varphi$. Then, Players 1, 3, and 4 have their goals satisfied and, therefore, do not have an incentive to deviate. On the contrary, Player 2 does not have its goal γ_2 satisfied. However, there is no beneficial deviation σ'_2 such that $(\vec{\sigma}_{-2}, \sigma'_2) \models \gamma_2$, otherwise, σ_1 would not be a winning strategy in the synthesis problem.

From right to left, assume there exists a strategy profile $\vec{\sigma}$ that is a Nash equilibrium. It is not hard to see that we have that $\vec{\sigma} \models \exists\varphi$, otherwise either Player 3 or Player 4 has a beneficial deviation. Moreover, the corresponding strategy σ_1 for Player 1 is winning for the system in the synthesis problem. Indeed, if by contradiction there exists a strategy σ'_2 for the environment in the synthesis problem, then we have that $(\vec{\sigma}_{-2}, \sigma'_2) \models \gamma_2$, and

so σ'_2 is a beneficial deviation for Player 2 in \mathcal{G}_φ , which contradicts the fact that $\vec{\sigma}$ is a Nash equilibrium.

For the RMG_F case, we can just regard the $\text{iBG}_F \mathcal{G}_\varphi$ as a RMG_F in which every agent i is associated with $|\Phi_i|$ modules, each of them allowing a single variable $x_i \in \Phi_i$ to be freely set at every iteration of the game execution. Thus, we obtain the assert. \square

Theorem 7 is a direct consequence of Lemma 2.

Games with Memoryless Strategies. In this subsection, we study games with memoryless strategies. We say that a strategy $\sigma_i = (S_i, s_i^0, \delta_i, \tau_i)$ for Player i is *memoryless* if $S_i = 2^\Phi$ and δ_i is deterministic. Intuitively, a strategy is memoryless if, for each state of the game, it always chooses the same action at such state. Moreover, a play $\pi^\infty \in (2^\Phi)^\omega$ is said to be memoryless if, for all $v, w \in 2^\Phi$, if $\pi_k^\infty = v$ and $\pi_{k+1}^\infty = w$, for some $k \in \mathbb{N}$, then, for all $h \in \mathbb{N}$, if $\pi_h^\infty = v$ then $\pi_{h+1}^\infty = w$. A profile $\vec{\sigma}$ made by memoryless strategies can only generate memoryless plays and vice-versa.

Moreover, it is not hard to build a polynomial size NBW $\mathcal{A}_{\text{memoryless}}$ accepting all and only the memoryless plays. This turns out to be useful in addressing the case of memoryless strategies. Indeed, to solve the NE NON-EMPTINESS problem with memoryless strategies, we only need to adjust the general procedure by pairing the automaton $\mathcal{A}_{\text{memoryless}}$ to every single component of the automaton $\mathcal{A}_{\text{NE}}(\mathcal{G})$. This operation then adds the memoryless requirement to the goal of a player and to the punishment strategies.

Now, although this solution technique allows one to prove that the set of Nash Equilibria in memoryless games is ω -regular, this is not optimal from a computational complexity point of view, which is still 2EXPTIME. For instance, a brute-force procedure can solve the problem in EXPSpace. Indeed, given the definition of strategies, we know that a memoryless strategy for a player in the game has (at most) 2^Φ states. Then, a memoryless strategy, as well as a strategy profile, can be guessed in time exponential in the size of Φ and saved using exponential space. In addition, using NE MEMBERSHIP we can check in PSPACE whether such a strategy profile is a Nash equilibrium of the game. Moreover, for some RMG_F the problem can be solved even with a better complexity. Indeed, if the overall number of guards in the modules is polynomial w.r.t the number of variables, we only have a polynomial number of states, and thus memoryless strategies can be polynomially represented, leading to a PSPACE complexity for the NE NON-EMPTINESS problem. Formally, we have:

Theorem 8. *The sets of memoryless Nash equilibria for iBG_F and RMG_F with both LDL_F and QPLDL_F are ω -regular. Moreover, the NE NON-EMPTINESS problem for iBG_F and RMG_F with both LDL_F and QPLDL_F with memoryless strategies can be solved in EXPSpace. For the case of RMG_F with an overall number of guards that is polynomial w.r.t. the number of variables, the problem can be solved in PSPACE.*

Games with Myopic Strategies. Another important game-theoretic setting is the one given by *myopic strategies* as they can be used to define *all* beneficial deviations. A game with myopic strategies is called a *myopic iBG_F*. We say that a strategy $\sigma_i = (S_i, s_i^0, \delta_i, \tau_i)$ for Player i is myopic if its transition function does not depend on the input variables, i.e., such that for each $s \in S_i$ and $v, v' \in 2^\Phi$, we have $\delta_i(s, v) = \delta_i(s, v')$. In a myopic

iBG_F, players are only allowed to use myopic strategies. In [?] it is shown how to reduce NE NON-EMPTYNESS for myopic iBG to the satisfiability of the QPTL formula

$$\varphi = \bigvee_{W \subseteq N} (\exists \Phi_1, \dots, \Phi_n. (\bigwedge_{i \in W} \gamma_i \wedge \bigwedge_{j \in N \setminus W} (\forall \Phi_j. \neg \gamma_j)))$$

where the formulae γ_i are the LTL goals of the players in the myopic iBG instance and the quantifier alternation is 1 (an alternation fragment for which the complexity is known to be EXPSPACE [?]).

To apply the solution provided in [?] to check the satisfiability of φ , one first has to transform each γ_i into the NBW automata recognizing their models. In the case of iBG_F, these LTL formulae are replaced by LDL_F formulae. However, as shown in the previous section, the infinite models of an LDL_F formula γ_i can also be recognized by NBW automata that are equivalent to some ω -regular expression of the form $\alpha \cdot (2^\Phi)^\omega$. Thus, in order to solve NE NON-EMPTYNESS for myopic iBG_F, we can first transform every LDL_F goal γ_i into the corresponding NBW \mathcal{A}_{γ_i} and then follow the technique used in [?]. Note that the same reasoning applies also for the case of QPLDL_F goals.

Moreover, for the case of RMG_F, we just have to replace the automaton \mathcal{A}_{γ_i} for γ_i with the automaton $\mathcal{A}_{\gamma_i} \otimes \mathcal{A}_G$, to force every player to comply with the modules specification.

We then obtain the following result for games with myopic strategies:

Theorem 9. *The NE NON-EMPTYNESS problem for myopic iBG_F and RMG_F with LDL_F or QPLDL_F goals can be solved in EXPSPACE.*

At this point it is important to note that a key observation behind this result is the fact that when playing with myopic strategies the strategies that are used to construct a run that is sustained by a Nash equilibrium (a Nash run) must be oblivious to players' deviations.

Games with Strong Nash equilibria. Despite being the most used solution concept in non-cooperative game theory [?], Nash equilibrium still has some limitations, for instance, it is not always stable and also it includes non desirable equilibria. As an example, consider a two-player game in which Player 1 controls a variable p and has the LDL_F goal $\gamma_1 = q$, while Player 2 controls a variable q and has the LDL_F goal $\gamma_2 = p^8$. It is clear that every strategy profile $\vec{\sigma}$ is a Nash equilibrium. Indeed, even in case a goal γ_i is not satisfied, the corresponding player cannot deviate from it, as the satisfaction of each player's goal is fully controlled by the other one. However, the desired outcome for both players is to satisfy both goals. Then, if we allow the two players to *collaboratively* deviate, the only stable outcomes are the ones making true both p and q at the first round of the computation.

A strong Nash equilibrium considers not only a single player's deviation, but also every possible coalition of players having a collective deviation incentive. Formally, for a given strategy profile $\vec{\sigma}$, we say that it is a *strong Nash equilibrium* if there is no subset

⁸Observe that γ_1 and γ_2 are propositional logic formulae, *i.e.*, special cases of LDL_F.

$C \subseteq \mathbb{N}$ and partial strategy profile $\vec{\sigma}'_C$ such that, for all $i \in C$, $\pi^\infty(\vec{\sigma}_{-C}, \vec{\sigma}'_C) \succ_i \pi^\infty(\vec{\sigma})$. Then, in a strong Nash equilibrium a coalition of players C has an incentive to deviate if and only if every player i in such a coalition has an incentive to deviate. By $\text{sNE}(\mathcal{G})$ we denote the set of strong Nash equilibria in \mathcal{G} . To check whether there exists a strong Nash equilibrium in a game using an automata-theoretic approach, we need to be able to express this notion of beneficial collective deviation with an appropriate automaton.

To do this, we just need to adjust the automaton $\overline{\mathcal{A}}_{\gamma_j}^{\mathcal{G}}$ given in the previous section, used to recognize all the plays that can be generated by a punishment strategy of the coalition \mathbb{N}_{-j} against j , having goal γ_j . Indeed, the concept of punishment can be easily lifted to punishing a group of players. To do this, for a set $C \subseteq \mathbb{N}$, consider the automaton $\mathcal{A}_C = \bigotimes_{j \in C} \mathcal{A}_{\gamma_j}^{\mathcal{G}}$ recognizing all the models that satisfy every γ_j , for $j \in C$. Then, as in the previous section, we can build the automaton $\overline{\mathcal{A}}'_C$ that recognizes the plays generated by a punishment strategy for the coalition $\mathbb{N} \setminus C$ against the goal being the conjunction of goals of coalition C . At this point, as in the case of Nash equilibrium, let us fix a set of “winners” $W \subseteq \mathbb{N}$ in the game and then consider the product automaton $\mathcal{A}_{\overline{2^L}} = \bigotimes_{C \in 2^L} \overline{\mathcal{A}}'_C$. By the semantics of the product operation we obtain that the automaton $\mathcal{A}_{\overline{2^L}}$ accepts those paths that are generated by some punishment profile, for each coalition of players $C \in 2^L$.

Thus, we have that the product automaton $\mathcal{A}_W \otimes \mathcal{A}_{\overline{2^L}}$ accepts exactly those paths for which every γ_i , with $i \in W$, is satisfied while, for each coalition $C \in 2^L$ the coalition \mathbb{N}_{-C} is using a punishment strategy against C . Now, as for Nash equilibria, we need to quantify over $W \subseteq \mathbb{N}$ to obtain an automata characterisation:

$$\mathcal{A}_{\text{sNE}}(\mathcal{G}) = \bigoplus_{W \subseteq \mathbb{N}} \mathcal{A}_W \otimes \mathcal{A}_{\overline{2^L}}.$$

The proof of correctness of this construction and its complexity is as for Theorem 4. Moreover, the same result can be obtained also with QPLDL_F objectives, as well as RMG_F games. Also, observe that the reduction from LDL_F synthesis provided for NE NON-EMPTYNESS with QPLDL_F goals can be reused with the same construction for the case of SNE NON-EMPTYNESS . Formally, we have the following result.

Theorem 10. *SNE NON-EMPTYNESS is in 2EXPTIME for both iBG_F and RMG_F , by using either LDL_F or QPLDL_F goals. In particular, for games with QPLDL_F goals, the problem is 2EXPTIME -complete. In addition, the set of strong Nash equilibria for games with either kind of goals is ω -regular.*

Remark 1. *The reader might notice that the automata product in the definition of $\mathcal{A}_{\text{sNE}}(\mathcal{G})$ contains a number of factors that is exponential in the number of agents. However, this blow-up does not affect the complexity of its emptiness problem. Indeed, every single factor is already of size double exponential. Therefore, a multiplication of an exponential number of doubly exponential sized automata is still of size doubly exponential.*

6. Concluding Remarks

Logic-based Multi-player Games Revisited. In the introduction section it was pointed out that the RMG/iBG and iBG_F frameworks rely on different automata techniques, and

that $\text{RMG}_F/\text{iBG}_F$ is better suited in certain scenarios. However, it is not the case that the $\text{RMG}_F/\text{iBG}_F$ frameworks generalise iBGs. Indeed, it should be noted that they are incomparable models. For instance, while iBG cannot be used to reason about games with goals over finite traces, $\text{RMG}_F/\text{iBG}_F$ cannot be used to reason about games with goals over infinite traces *only*, that is, regardless of the satisfaction of players' goals in the associated finite traces; from a logical point of view, while RMG/iBG considers LTL, $\text{RMG}_F/\text{iBG}_F$ can handle goals in LDL_F , which on finite traces is strictly more expressive than LTL_F [?], and also than LTL over finite traces. However, as shown here using new automata techniques, the complexities of some problems in each game model coincide in the worst case for many variants of these different kinds of games.

Automata for Linear Dynamic Logic. Logic, games, and automata are intimately related; see, *e.g.*, [? ?] and references therein for examples. We build upon the automata constructions for Linear Dynamic Logic (LDL_F [? ?]), which were introduced to solve the satisfiability and synthesis problems for LDL_F over finite traces. Specifically, we have initially used such constructions to translate LDL_F formulae to alternating automata on finite words (AFW) and, based on them, we have defined *new and optimal automata constructions* that characterise the existence of (strong) Nash equilibria on top of the standard Boolean games framework.

However, even though most, but not all, of the automata constructions we have presented in this paper are optimal, they still enjoy two useful properties. Firstly, that they are strongly based on automata on finite words, with only an extension to deal with infinite runs, a feature that could be used a lot further. Secondly, that such automata constructions recognise the sets of Nash runs, which, as shown in this paper, makes them extremely useful from a semantic point of view. Indeed, using other automata approaches, *e.g.*, for iBGs, our expressiveness results do not easily follow.

In addition, the automata constructions defined in this paper can be modified to reason about other game settings, making it a rather widely adaptable reasoning technique. For instance, we believe it is also possible to extend *some* of the results we have obtained to two-player games with imperfect information. This should be possible, in some cases, using recent automata constructions to reason about LDL_F formulae under partial observation [?].

Imperfect Information. Following the research line delineated in [? ? ?], one might wonder about the complexity of solving $\text{RMG}_F/\text{iBG}_F$ in the context of imperfect information. It is important to notice that the synthesis problems for LTL for both perfect and imperfect information is decidable [? ?]. On the other hand, the NE-NONEMPTINESS problems for games having LTL goals is decidable for the perfect information case [?], but undecidable for the imperfect information case [?]. Similarly, regarding LDL_F goals, the synthesis problem is decidable for both perfect [?] and imperfect [?] information, while we here prove that the NE-NONEMPTINESS problem is decidable. This suggests that the same problem might be undecidable under the imperfect information assumption, as it is for conventional iBGs. However, as the expressive power of LDL_F is incomparable with the one of LTL, it is not clear whether the undecidability proof (which strongly relies on the expressiveness of LTL) can be retained in this case. Moreover, it has been shown that for specific cases of imperfect information in games

with LTL objectives, the problem might be decidable [? ?]. For this reason, we plan to address this question in future work.

Verification and Equilibria in Logical Form. Here we have addressed the “rational verification” problem [?] of multi-agent systems using an automata-theoretic approach. We believe that these automata constructions may be used as the underlying models of a variant of strategy logic [?] (SL) based on LDL_F over finite traces. We will do so in future work. SL is not the only option to investigate. Many other logics for strategic reasoning can be found in the literature; see, *e.g.*, [?]. We believe this is a promising research direction, which would allow us to reason about the behaviour of games over finite traces using other solution concepts in a uniform logical framework. A step toward this direction has been recently taken by the formal methods community [? ?].

Expressiveness and Automata Characterisations. Nash equilibrium, and other solution concepts, have been characterised in other works using a range of automata and logical approaches. See, for instance, [? ? ? ?] for some examples. Although delivering optimal constructions, it is not clear that using such techniques one can characterise the sets of Nash runs in a game. A key feature and contribution of our automata constructions is that they can be used to characterise such sets of equilibrium runs in a uniform way.

Acknowledgments

This paper is an extended version of [?]. We acknowledge with gratitude the financial support of the ERC Advanced Investigator grant 291528 (“RACE”) at Oxford.

References